The Prisoner's Dilemma and Laissez Faire

Revision 9/25/2009

The Dilemma

You're walking down the street, minding your own business, when a van screeches to a halt besides you, half a dozen guys pile out, throw a black bag over your head, muscle you into the back of the back of the van, and stun you unconscious. You wake up in a windowless room with a one-way mirror, sitting in a chair, handcuffed, with some guy shoving you and telling you to wake up.

He says his name is George. He flashes some badge too fast for you to read it. He says you're in some secret black-site prison in a third world country that's too difficult for him to pronouce. He tells you some warlord handed him another prisoner, whom he only refers to as "Alice", for a $10,000 bounty. The warloard said she was on the battlefield when he caught her. Then he tells you your options:

You can rat on Alice and testify before a secret military commission that Alice committed terrorism or you can clam up, not say anything about Alice, and claim your innnocence.

He tells you that he's making the same offer to Alice, offering her the chance to rat on you or clam up.

Tomorrow, he'll take you before a three judge military panel and you can choose to rat on Alice or clam up. Then he'll bring you back to a holding cell. Then he'll take Alice to the judge, and she can either rat on you or clam up. Then he'll take her back to a holding cell. Neither one of you will know what the other is choosing until after you're back in your holding cell, until after you've made your choice before the judges panel.

The three-judge panel decides in secret and keeps no records of evidence or witnesses. They will simply hand down their sentences. There are three possible outcomes:

> (A) If one of you rats and the other clams up, the rat will go free that day and the clam will be sentenced to 10 years for acts of terrorism.
>
> (B) If you both rat, then you'll both be sentenced to 5 years.
>
> (C) If you both clam up, then you'll both be held for 1 year and released without charges.

You ask him why he's doing this. George gives you a blank stare for a moment, then he tells you about the "private contractor of the month" award that comes with three days vacation. All he has to do is bring enough people before the judges and he'll win.

You protest something about "rights". He hits you with a stun gun and you black out. You wake up in a prison cell alone. You've got one day to decide whether to rat on Alice or clam up.

Welcome to the Prisoner's Dilemma. It's a game theory concept that Merrill Flood and Melvin Dresher working at RAND came up with back in 1950. It's often presented with a generic narrative form, usually involving two bank robbers. But I decided to update it slightly.

The Prisoner's Point of View

So, what do you, as Bob, do? Do you rat out Alice or clam up? Alice can rat or clam up to. That leads to four possible outcomes. Game theory uses a "payoff matrix" to represent the choices:

		ALICE			
		rat		clam	
BOB	rat		5		10
		5		0	
	clam		0		1
		10		1	

The numbers in the columns represent the number of years each prisoner gets.

There are four possible outcomes:

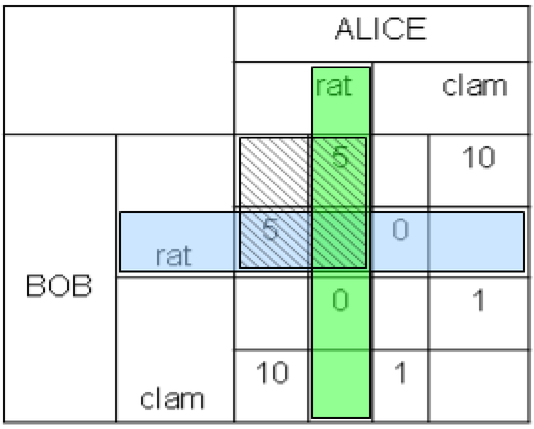

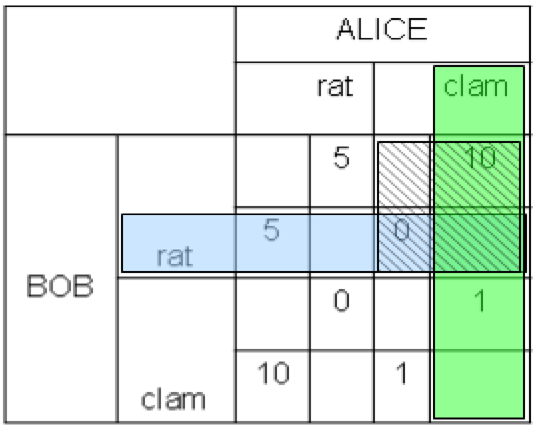

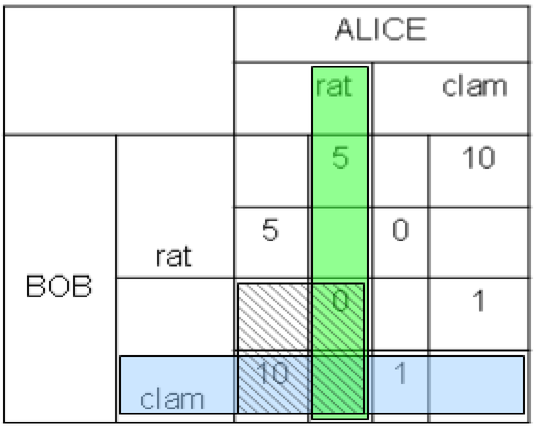

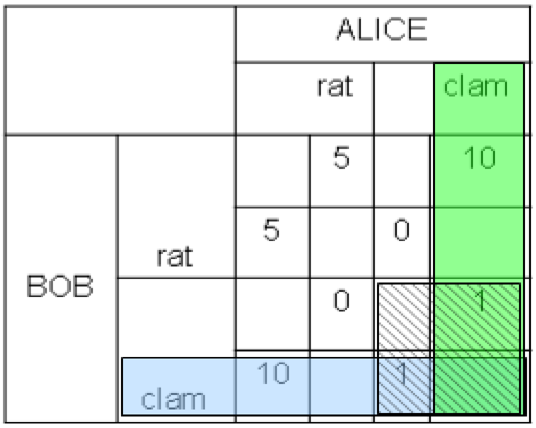

From your point of view as Bob, you've got two options:

Rat: resulting in 0 or 5 years in prison

Clam: Resulting in 1 or 10 years in prison

If your motivation as Bob is to minimize your prison sentence, then you would select Rat, minimizing your prison time to, at most, 5 years. If you clam up, you risk getting 10 years in prison.You have no idea who Alice is or what her motivation is. You have no idea what she might choose to do. Therefore, to minimize your prison time, you have to choose the option which has the lowest possible prison times.

When you, Bob, find yourself in a Prisoner's Dilemma as described above, and you want to minimize your time in prison, you should always Rat on Alice.

The Dilemma about the Prisoner's Dilemma

If Bob finds himself in a Prisoner's Dilemma, he realizes that to minimize his time in prison, he must choose to Rat on Alice. This minimizes his time in prison to 0 or 5 years. But then that would mean that if Alice is presented with the same dilemma, she too must choose to Rat on Bob. This means that if Alice and Bob find themselves in a Prisoner's Dilemma situation, they would both choose to rat on the other, and they would both get 5 years.

This results in a non-intuitive outcome: Two people both acting in their own best self interest achieve the second worst possible outcome.

Out of the 4 possible sentences: 0, 1, 5, and 10 years, both prisoners get 5 years.

Fallacies about the Dilemma

This dual-rat outcome is sufficiently counter-intuitive for some that they disagree with the result. Some will assert that both Alice and Bob would see their options and both would clam up, resulting in both people getting a 1 year prison term.

Because the Prisoner's Dilemma is used as a model for some very important, real-world problems, if a person arrives at the wrong conclusion in the model, they'll tend to arrive at the wrong conclusion in the real world. So, it's important to understand the model and the correct answer.

There are several possible misunderstandings that might occur and cause someone to arrive at the incorrect conclusion that Alice and Bob would both clam up. I'll try to cover some of the most common erroroneous paths here:

Fallacy of Omnicient Viewpoint

The Prisoner's Dilemma is intended to be taken from the prisoner's point of view. You are presented the Prisoner's Dilemma as if you were one of the prisoners. In the examples above, you are cast into Bob's point of view. You know nothing about Alice. You have to decide what is best for you. You have to decide what is best for Bob.

If someone approaches the Prisoner's Dilemma from a "meta" point of view, they might reach a different conclusion. Unfortunately, in its classical form, the prisoner's dilemma ("PD") is presented from a "meta" view:

> Two suspects are arrested by the police. The police have insufficient evidence for a conviction, and, having separated both prisoners, visit each of them to offer the same deal. If one testifies (defects from the other) for the prosecution against the other and the other remains silent (cooperates with the other), the betrayer goes free and the silent accomplice receives the full 10-year sentence. If both remain silent, both prisoners are sentenced to only six months in jail for a minor charge. If each betrays the other, each receives a five-year sentence. Each prisoner must choose to betray the other or to remain silent. Each one is assured that the other would not know about the betrayal before the end of the investigation. How should the prisoners act?

Notice that this narrative is different from my version at the beginning of this document. It's not just that I switched from bank robbers to terrorists, it's the context of the final question. In my version, the question is this: What should you, as Bob, do? In the classic form, the question is: What should both prisoners do?

Most novels are told from a third-person-limited point of view. "Third person" means the characters are refered to using third person pronouns, such as "he" and "she". The "limited" means that the narration is limited to following one charater. A problem with the "classic" form of the Prisoner's Dilemma is that it is told from a third-person-omniscient point of view. The "omniscient" bit means that the narator can see anything going on at any time through a god-like, all-knowing perspective. The narrator can inform the reader what all the characters (Alice and Bob) are doing and thinking at a particular time.

In a mathematical sense, the point of view of the narrative is irrelevant. Game theory experts understand that the question is really intended from an individual prisoner's point of view. But if a game theory hobbyist (or some layman off the street) takes the question "What should both prisoners do?" and answers it in a literal interpretation, an omniscient interpretation, they get a fundamentally different answer.

So, you're some god-like being who commands both Alice and Bob. They're in the prisoner's dilemma. What should they do?

You, as god, realize that if both prisoners clam up, then both will get 1 year in prison. That's a total cumulative time of 2 years. The alternative is one rat and one clam, which gives 5 for Alice and 0 for Bob, (or swap the names) for a total cumulative time of 5 years. So, mathematically speaking, you choose to have both Alice and Bob clam up, get 1 year apiece, and a total time of 2 years.

So you as omniscient god hand down the sacred script that tells your actors, prisoners, subjects, what to do. And what you tell them is "Thou shalt clam up when faced with the prisoner's dilemma."

Which is fine when both prisoners facing the prisoner's dilemma are your people, your subjects, your mortals. Your subjects know the commandment is to clam up, so they clam up.

But problems arise when the other prisoner follows a different God, or even an atheist who is simply following the mathematically best solution for HIS point of view.

If you are Bob and are trying to do what is best for you, you would choose to rat on the other prisoner. If you are God and are trying to do what's best for both prisoners, you would command them to clam up. So, depending on the narrative point of view, omniscient or limited, you get a different answer.

Fallacy of Preagreement

You, as Bob, never met Alice before. You are only told that she exists. Unfortunately, the "classical form" of the prisoner's dilemma confuses the reader about this too. In the classic form, the prisoners are two bank robbers. The narrative goes on to say that the police have "separated" the two robbers, which suggests that they were together at some point. Possibly working together. Possibly working together to rob the very bank that the police are investigating. If so, then it wouldn't be too much of a leap for a person to realize that the best solution is for the robbers to agree ahead of time that if the job goes bad to clam up and get a lawyer.

That's not what the actual prisoner's dilemma allows for. Not in the classic sense at least. You have never met Alice. And you only know she exists because you are told she exists.

Fallacy of Post-Release Enforcement

This is an extension of the preagreement fallacy. If the two robbers agree beforehand to clam up if they're nabbed by the cops, there is really nothing that forces either one of them to comply with that agreement. If Alice rats on Bob and Bob clams up, Alice will go free immediately. Bob will have to wait 10 years to get out of prison before he can find Alice and extract some payback from her in some form.

Alternatively, the robbers might involve a third party, such as a mob boss who is directing them to rob the bank, and who directs them to clam up if picked up by the cops. And if the mob boss sees Alice get out of prison immediately, he might have her picked up by one of his men and hold her until Bob is released too. If Bob isn't released, the boss can punish Alice for betraying him.

If you incorrectly assume that the robbers can put some sort of post-release enforcement policy in place, then you would incorrectly arrive at the idea that Alice and Bob would both naturally clam up. But the original prisoner's dilemma does not allow for this. You, as Bob, don't know who Alice is, don't know what she will do, and will never see her again in your life.

Fallacy of Third Party Intermediators

Another way one might alter the prisoner's dilemma is to involve some third party, Trent. Trent is someone you know, usually someone you trust. And you might use Trent to threaten Alice or to spy on Alice and try and figure out who she is and what sort of choice she might make. Maybe you use Trent to negotiate some sort of arrangement with Alice so that you both clam up.

The "omniscient god" and the "mob boss" are types of third party intermediators.

These are not the prisoner's dilemma games anymore. They are different games with different tables of choices and with different outcomes.

Fallacy of "You Think You Know Alice"

Let's say that for some reason, you think you know something about Alice. You think Alice is really an innocent person like you who was picked up by George and his secret prison system. You decide that Alice is innocent and you're innocent, so you figure she wouldn't want to rat on you just like you really don't want to rat on her. Therefore, you decide that the best thing for you, as Bob, to do is to clam up.

Well, Bob, I told you that you don't know who Alice is. But you didn't listen. After you clam up, it turns out that "Alice" was George's nickname for Al Asad, a ruthless terrorist who executed Al Funani and is leader of a group of ultra nationalists. You are sentenced to 10 years. You die of hypothermia 7 years later and are buried in an unmarked grave behind the chemical shed.

It is true that if you know who Alice is, or more specifically, if you know what choice she will make, then you can make the best choice for yourself. If she clams up, you can either be nice, clam up too, and get out in a year, or you can rat her out and go free immediately. If she rats you out, you will probably want to rat her out too so that you only stay in prison for 5 years instead of 10. But knowing that she ratted you out, you may feel less guilty about ratting her out as well.

The thing is that you don't know anything about Alice. Nothing. Nada. Zip.

In the 1980's, Douglas Hofstadter wrote "Metamagical Themas" which discussed, among other things, the idea of "Superrationality", which is a concept he describes around the prisoner's dilemma. Superrationality suffers from the fallacy of knowing who the other prisoner is. If you know the other prisoner is "superrational", then you too can follow a superrational strategy.

But if you, Bob, are superrational based on the condition that Alice is superrational, then from Alice's point of view, you are not superrational. You're conditionally superrational. Which means Alice can act non-superrationally if she needs to as well. Which means no one is superrational unless they already know who the other prisoner is.

And that's not the same game as the prisoner's dilemma.

Fallacy of a Sequential Game

You, as Bob, must choose without knowing anything about Alice or her choice. Alice must choose without knowing anything about Bob or his choice. Effectively both choices happen simultaneously.

If Alice chooses first and she chooses to clam up, and you got to see her choice, then you would see that she sacrificed the possibility of getting out immediately for getting out, at the earliest, in 1 year. You then might feel motivated to clam up as well, so you both get out in 1 year. But that's a sequential game, not the prisoner's dilemma.

Superrationality could also be looked at as a sequential game. You are superrational on the condition that Alice is superrational. The only way to know if Alice is superrational is to know her choice. If you know her choice, you're playing a sequential game.

Fallacy of Implied Iterations

The notion of the prisoner's dilemma is that it is a one-shot deal. You, as Bob, will never see Alice. And after you've served any sentence, you won't see Alice or George every again. In the "classic" narrative, two bank robbers are picked up by the police. One could assume that there is some probability that the bank robbers will end up seeing the police again in the future. At which point, multiple instances of the prisoner's dilemma played in sequence becomes an "iterated" prisoner's dilemma. The same two prisoners keep facing the police over and over.

The thing about an iterated prisoner's dilemma is that it is a different game than a one-time prisoner's dilemma. In an iterated prisoner's dilemma, you're Bob, and every once in a while, you get picked up by the cops, and they always pick you up with your known associate, Alice.

When you, as Bob, are faced wth an iterated prisoner's dilemma, you are faced with a different game. In a one-time prisoner's dilemma, there is no communication channel between the two prisoners. And there is no post-release enforcement of any kind, either by you or some third party.

But with the iterated prisoner's dilemma, you remember how Alice chose in previous rounds. And her previous choices tells you something about Alice. They can become a communications channel of sorts to let you know a little bit about Alice. Also, an iterated prisoner's dilemma has the implied notion of future dilemmas. You can use future dilemmas as a promise to reward or punish the other prisoner.

This results in a different game. A one-time prisoner's dilemma plays like a game of paper-rock-scissors. You make one choice and on the count of three, you reveal your choice. End of game. An iterated prisoner's dilemma is more like a game of checkers. You make individual moves over time, and those individual moves taken as part of a longer sequence results in a single "game" of iterated prisoner's dilemma.

The result is that you can get cooperation between prisoners. Both can clam up, knowing that in the next iteration the other prisoner can reward or punish them. In an iterated prisoner's dilemma, it can make mathematical sense for prisoners to cooperate.

But that isn't a one-shot prisoner's dilemma. It's a different game with a different outcome. If you approach the one-shot prisoner's dilemma as if it is the first of many dilemmas between you and another prisoner, then you're committing the fallacy of implied iterations.

One way to see if you're committing the fallacy of implied iterations is to alter the payouts of the game so that you are more consciously aware of the fact that there is only one iteration and it will never happen again. If Alice rats on you and you clam up, instead of getting 10 years in prison, you are taken out to the chemical shed and executed.

If that gives you a different answer, then it may be because you're implying iterations when there are none.

Fallacy of Moral Strategy

The fallacy of morality is based on a human attitude that can emerge because the circumstances of the prisoner's dilemma are imposed against someone's will and they may be punished based on what action you choose. You, as Bob, may feel bad about ratting on Alice because if Alice didn't do anything wrong, then your action will put her in jail for 10 years.

One way to see if you're committing the fallacy of morality is to alter the narrative so that both Alice and Bob chose to enter into the circumstances. If that causes you, as Bob, to chose a different action, then you may be bringing some sense of morality into the situation.

In the prisoner's dilemma, you, acting as Prisoner Bob, can choose to act however you wish to act and for whatever motivation you wish to act for. However, the prisoner you are playing against may be acting in her own selfish interests, and may have no qualms about letting you sit in jail for 10 years.

Many react to the idea of "ratting" out on someone to be an immoral act, and so they alter their choice based on the moral notion that "ratting" is bad, so they choose to clam up. This is another way in which morality can enter into the narrative and affect it in unintended ways.

A Non Iterated Amoral Prisoner's Dilemma

To see if you are approaching the prisoner's dilemma using an implied iterations fallacy, a fallacy of a moral strategy, or other fallacy, here is the exact same dilemma from a game theory standpoint, but presented in a narrative that is clearly a one-time game and is much less likely to trigger subconscious moral strategies.

The Millionth Customer Dilemma

You're Bob Smith. You walk into the World Famous Grand Hotel to check in for the night, and all the lights start flashing, bells are going off, and confetti starts falling from the ceiling. The manager rushes out to you and says "Congratulations, you are our one-millionth customer!" He then explains to you that you have a chance to win a prize. He whisks you off into a nearby office and hands you a form to fill out.

At the bottom of the form, you're a bit confused. It asks you whether you want a king sized bed or a queen sized bed, and the form says your choice will affect what prize you get. The manager explains that the millionth-and-one customer, Alice, is in another office filling out the same form. Then he tells you how the prize is calculated.

> If you choose King and Alice chooses Queen, you get 20 million dollars and Alice gets 1 million dollars.
>
> If you choose Queen and Alice chooses King, you get 1 million and Alice gets 20 million.
>
> If you both choose King, then you both get 5 million dollars.
>
> If you both choose Queen, then you both get 15 million dollars.

The manager tells you how happy he is that they were able to get to the one-millionth customer before next month. Next month, he explains, the hotel is being demolished by new owners who are putting up a sports arena. The original hotel owners are getting out of the business entirely and retiring to the Caribean.

You didn't see who was in line behind you, so you have no idea who Alice might be. You take out your cell phone, but the manager informs you that they're in a Faraday cage to prevent cheating. And the door is locked from the outside. You cannot leave the room until you've selected King or Queen.

So, what do you, as Bob, do?

A reader might find this narrative much less emotionally charged. There is no "ratting out" your bank robbing buddy Alice in this narrative. There is no sending someone to prison or to some horrible fate that might happen behind the chemical shed. No matter what, you and Alice both get a nice prize that neither one of you were getting before you walked into the hotel.

Since the hotel is being torn down and the owners are retiring, there's absolutely no chance that you and Alice will find yourselves the millionth and millionth-and-one custermers of the hotel again. It is far less likelyhood that people will commit the fallacy of implied iterations with this narrative.

You still have two options, they're just labeled emotionally neutral terms of king or queen sized beds. And there are still four possible payouts, but measured in millions of dollars you and Alice win, rather than years you and Alice spend in prison. The prizes are 1, 5, 15, and 20 million dollars.

While some readers would feel bad about sending Alice to prison for 10 years because they chose to rat her out, I believe fewer readers would have quite such negative feelings about choosing an option that caused Alice to win "only" 1 million dollars for a chance for them (Bob) to win 20 million dollars.

The four outcomes look like this:

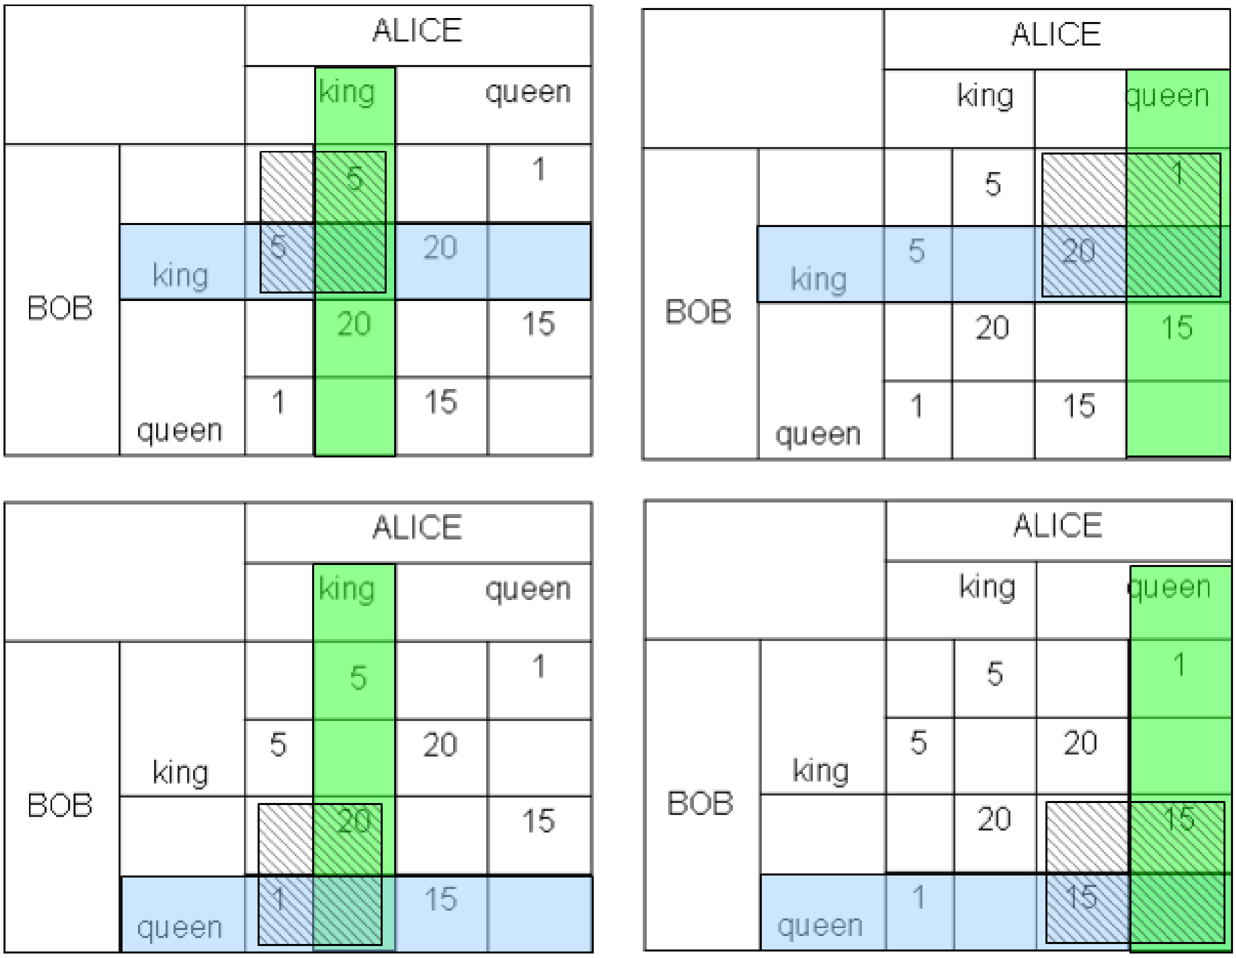

You, as Bob, have two choices. If you choose the King sized bed, you may win 5 or 20 million dollars. If you choose the Queen sized bed, you may win 1 or 15 million dollars. If you choose King, you're guaranteed at least 5 million dollars. If you choose Queen, you might only win 1 million dollars. So if you want to maximize your possible winnings, you would select King.

But the millionth-and-one customer, Alice, is presented with the same scenario. So, she'd choose King as well. But both players choosing King means both players win "only" 5 million dollars. If both had chosen Queen, then both would win 15 million.

Once again, we see the same dilemma. Two people are put in a situation where they don't know the other person, don't know what the other person will do, and have no way to communicate with the other person. Neither can make any sort of pre-agreement. Neither can effect any sort of post-winning punishment that would enforce the other person to choose one size bed over another. There are no iterations whereby Alice and Bob might eventually see a pattern in how the other person votes or use their current vote to reward or punish previous votes.

And if both are motivated to maximize their individual winnings, then both choose King, and both end up getting the third smallest prize of 5 million dollars (out of a possible 20, 15, 5, and 1 million dollar prizes)

If both take a chance that the other person goes for a queen, then they both get 15 million dollars. But if Bob takes the chance and chooses Queen, Alice might choose King and Alice would get 20 million dollar payout, leaving Bob with only 1 million. Alice doesn't know that Bob chose Queen. She must make her choice without knowing Bob's choice.

I believe my narrative of the Millionth Customer Dilemma is more likely than the classic form of the prisoner's dilemma to be read by non-experts of game theory and have those people understand the dilemma itself without them choosing a path for moral or emotional reasons. I believe that the Millionth Customer Dilemma also uses a narrative that avoids problems like the omniscience fallacy or the implied iteration fallacy and so on.

I believe more people would read the Millionth Customer Dilemma and see that "King" (for 5 or 20 million dollars) is better than risking "Queen" (for 1 or 15 million dollars), than people would read the classic form of the prisoner's dilemma and understand that your best option, as Bob, is to rat on Alice.

Randomizing Alice

With the emotional and moral issues attached to going to prison and sending someone else to prison removed from the narrative, some may find it more natural to look at the millionth customer dilemma as a gamble.

If you, as Bob, don't know anything about Alice, then Alice's move can be viewed as a single random event which you have to deal with. Alice has two choices, king or queen. Alice is random. Therefore Alice's choice could be like a coin toss. Heads she picks king. Tails she picks queen.

We can remove Alice completely from the narrative by replacing her with a coin toss performed by the hotel manager.

If you, Bob, pick king, then you've got a 50% chance of getting 5 million dollars and a 50% chance of getting 20 million dollars. If you choose queen, then you've got a 50% chance of getting 1 million dollars and a 50% chance of getting 15 million dollars.

King => 50% chance of $5m. 50% chance of $20m.

Queen => 50% chance of $1m. 50% chance of $15m.

Given the above choice, it might become even easier for people to see that their best chances to maximize their winnings is to select king.

Randomizing Alice may also help some people understand that the prisoner's dilemma consists of two people making independent decisions. Alice doesn't know how Bob will choose. Bob doesn't know how Alice will choose. You, as Bob, must choose king or queen based on some event that you cannot predict: a coin toss.

If you're committing the fallacy of thinking you know Alice, then you should be able to clearly see this fallacy once you replace Alice with a coin toss. You don't know anything about Alice. Her choice becomes a random coin toss.

Some may argue that the population from which Alice is taken must have some distribution of behaviour, therefore you should be able to predict Alice's choice based on a statistical analysis of the population. OK. Fine.

Wikipedia mentions a study that found that when presented with a straightforward prisoner's dilemma about 40% of the people choose to clam up. 60% ratted out the other prisoner. Using the milltionth customer dilemma, that means the "coin" is slightly lopsided. 40% would choose queen. 60% would choose king. Your odds are recalculated here:

King => 60% chance of $5m. 40% chance of $20m.

Queen => 60% chance of $1m. 40% chance of $15m.

If you're trying to maximize your winnings, you, as Bob, would still select King.

Now, if the population distribution swung the other way, if population studies found that 40% would rat and 60% would clam up (40% would choose king, 60% would choose queen), then Bob's choice looks like this:

King => 40% chance of $5m. 60% chance of $20m.

Queen => 40% chance of $1m. 60% chance of $15m.

And again, you, as Bob, would want to select king to maximize your winnings.

Empathizing with a Coin

Some might feel badly about sending Alice to prison for 10 years and will tend to avoid ratting her out. Replacing Alice with a coin will remove the fallacy that comes from implementing a moral strategy.

Forcing a Non Sequential Game with a Coin Toss

A coin toss has the same probability of outcome regardless of how you choose to act. Whether you choose king or queen, the coin toss is still 50/50 or whatever. The results do not change because of the choice you made. Replacing Alice with a coin toss may help some see that they're committing the fallacy of a sequential game, thinking Alice will choose based on your choice.

The Fair Winnings Strategy

Another result of empathy might be to use a fair winnings strategy. You discard any result in the payoff matrix that produces extremely lopsided winnings, something deemed "unfair". You then select among the remaining "fair" outcomes. In the millionth customer dilemma, that would mean the 1 million for one player and 20 million for the other player would be ruled "unfair" and those options would be discarded out. You, as Bob, are then left with two options, king or queen, but with different outcomes.

King => $5m fair winnings

Queen => $15m fair winnings

In this case, you, as Bob, using the fair winnings strategy would choose queen, hoping both win $15 million dollars.

In the prisoner's dilemma, the "fair" outcomes are 1 year each and 5 years each.

Clam up => 1 year

Rat out => 5 years

In this case, you, as Bob, using the fair winnings strategy would choose to clam up, hoping that both of you are sentenced to only 1 year.

The problem, again, is that you, as Bob, can only rule out the "unfair" results if you know Alice is also fair in her strategy. And you don't know that. This is really a subset of the "you think you know Alice" fallacy. It's like the notion of superrational players. You, as Bob, are really only conditionally superrational, if you know Alice is superrational. Likewise, you can only be fair if you know Alice is fair as well. You are conditionally fair.

If you choose based on a "fair winnings" approach, but Alice chooses based simply on maximizing her potential winnings, then you'll only get 1 million and Alice will get 20 million. Just because you implement a fair winnings strategy doesn't remove the possibility that Alice might choose something you deem "unfair". Choosing based on a "fair winnings" strategy doesn't actually alter the possible outcomes, the matrix, of the game.

Player Strategies Do Not Alter the Payoff Matrix

A player's strategy for selecting their choice in the dilemma does not alter the payoff matrix of possible results.

Some people think if they choose what is "fair" that Alice would or should choose what is "fair" too. Regardless of your choice, Alice still has all the options available to her in the payoff matrix and she can choose based of her own strategy.

A Con Man Playing in a Sea of Fair Players

Let's say that you happen to know that the vast majority (say 95%) of people in the entire population use a fair winnings strategy. But lets say that you yourself are a con man. You know you're playing a one-time game. You know you'll never see the other player ever again. You know the other player cannot hunt you down and punish you in any way. You know the person is going to play "fair" expecting you to play "fair", but you don't care. You want the best outcome for you, and you're not afraid to play on the other person's sense of fairness to get it.What would you, Charlie the Conman, do in the millionth customer dilemma?

If you know the other player is chosen from a population of people where 95% of those people would chose the fair option, i.e. Queen, then your choices look like this:

King => 5% chance of $5m. 95% chance of $20m.

Queen => 5% chance of $1m. 95% chance of $15m.

Compare this to the conman playing in a game where he knows nothing about the other player. Charlie must revert back to the other player as if they were a 50/50 coin toss.

King => 50% chance of $5m. 50% chance of $20m.

Queen => 50% chance of $1m. 50% chance of $15m.

This suggests that the more the population tends towards being fair, the more a conman can take advantage of the people in a one-time game. The more payoff per game for the conman.

The more fair the population is, the more the population as a whole benefits. Most outcomes will see both players get $15 million each. But the more fair the whole population is, the more easily a single conman can make large payoffs from individuals in the population, because those individuals will tend to trust the other player to play fairly.

When one player is playing a "fair winnings" strategy and the other is playing only for their own benefit, the "fair" player gets the worst possible outcome. This is sometimes called the "Sucker's Payoff" in prisoner dilemma discussions.

Summary of the One-Time Prisoner's Dilemma

Alice and Bob know nothing about each other. Both must choose simultaneously without seeing the other person's choice. Both want to maximize their personal benefit, but because they have no knowledge of each other, and due to the circumstances of the game, they both end up with the second worst possible outcome.

All the fallacies around the prisoner's dilemma boil down to the player believing they know something that they really don't know. They think they know how the other player is going to choose.

This may be because they incorrectly think the game is iterated when it isn't. They may think there is some kind of preagreement when there isn't. They may think there is some kind of post-game enforcement when there isn't. They may view the game from an omniscient point of view when it isn't. They may believe that the game is sequential, that they choose, then Alice sees their choice, and then Alice's chooses, when it doesn't at all work that way. They may think they can somehow use a fair winnings strategy to alter the payout matrix, but they can't.

It can be difficult for people to hold the notion of "I don't know" and not accidentally fill in that void with false knowledge. They think they know, when they don't. They reach a result based on that false knowledge and it's the wrong result because the false knowledge is, well, false.

You, as Bob, DO NOT KNOW how Alice will choose. Alice DOES NOT KNOW how you, Bob, will choose. You both must make your choices WITHOUT KNOWING how the other will choose. That the players must make their choices WITHOUT KNOWING all the information is what results in both players getting the second worst outcome.

It is this missing knowledge that creates the dilemma.

But filling in this knowlege with false information only yields even worse outcomes, the sucker's payoff for some players.

One of the fundamental lessons to learn from the prisoner's dilemma is that we don't know everything about the person we're dealing with and making decisions with that information missing yeilds less than ideal results. And as much as we'd like to improve our results, it is a fallacy to invent non-existent information to say the outcome should be better than it really is.

You Don't Know That You Don't Know, or Dogmatic Strategy

A dogma is an authoritive belief that is not to be disputed or diverted from. At the core of the one-time prisoner's dilemma is the fact that there are some things you don't know, and there are important decisions you must make without complete information.

Filling in this unknown with false information is to create a dogma. How you choose your actions in the prisoner's dilemma is your strategy. Choosing based on false information that you don't really know is using a dogmatic strategy.

In a one-time prisoner's dilemma, where you don't know anything about the other player or how they will choose, you must rat on the other prisoner if you want to minimize your time in prison. If you assert that you know anything about the other player, that to minimize your time in prison you should clam up, then you are using a dogmatic strategy. You think you know something about the other prisoner that you don't. The other prisoner is like a coin flip to you. You can't affect it and you don't know which way it will land.

The thing about dogma is that people often cling dearly to them and resist attempts by others to show them the error of their ways. If you are committing some dogmatic strategy around the one-time prisoner's dilemma, then I'm not sure what to tell you at this point. I'm not sure any explanation will change your mind.

If you think there is a better solution to the one-time prisoner's dilemma than ratting out the other prisoner, then the first thing I'd say is to make sure you read this entire document from the beginning without skipping anything. If you've read everything this far, and still think you've got a better strategy, then I'm out of explanations for you.

The Iterated Prisoner's Dilemma

The iterated prisoner's dilemma has Alice and Bob choosing to rat or clam on each other over and over again. The one-time prisoner's dilemma is a single game unto itself. It's like rock-paper-sciccors; you make your choice, you and the other player reveal your choices simultaneously, and the outcome is determined.

The iterated prisoner's dilemma is different in that each choice is a single turn in a longer, sequential game. The iterated prisoner's dilemma is sort of like a game of checkers. Moving a piece on the board is like each player making a single choice to rat or clam up. The running total of those choices gives you the final tally of the entire game.

In game theory, when analyzing an iterated game, instead of a payoff matrix, you have to use something called an extensive form or a "tree". The "branches" of the tree represent each player's decisions. Each branch comes to a "node". Each "node" represents a new choice the players have to make. Along each node, a running tally can be made of each player's score. And at the terminal nodes, the game ends and each players final score is determined. It's rather complicated and I'll leave it as an exercise for the reader to look up the details. Suffice it to say, that in an iterated prisoner's dilemma, you can't look at a payoff matrix of one turn to determine the entire game. You have to look at the extensive form, the tree, the running tally of all the choices.

The important difference about the iterated prisoner's dilemma versus the one-time prisoner's dilemma is that in the iterated game, some limited communication is possible and some post-agreement enforcement is possible. These differences yeild a different outcome.

In a one-time game, both players seeking their own self interest would vote to rat out the other player. This yeilds the second worst outcome for both players.

Once you start iterating the game, though, Bob might sacrifice himself, clam up, and chance getting the absolute worst possible outcome. Why? Because it can be used as a signal to Alice that Bob is willing to cooperate. It communicates, essentially: *"Look, I know my best bet is to rat on you. But I'm willing to cooperate with you so that we both get the best "fair" outcome of 1 year."* If Alice manages to extract that meaning from the single action of Bob clamming up, in the next iteration, Alice realizes that if she clams up too, they both get 1 year instead of the 5 years they're getting now.

But the sucker payoff remains. If Alice continues ratting out Bob, and Bob tries clamming up, each time that happens, Alice goes free immediately.

But if Alice continues to rat on Bob, then Bob can respond by changing from clamming up back to ratting on Alice. At which point Alice goes from getting out immediately to getting out in 5 years again. Bob can use the iterative aspect of the game to "punish" Alice if she doesn't cooperate.

The iterative prisoner's dilemma has a very limited communications channel and a very limited post-choice enforcement mechanism.

And that channel can be abused. For example, Bob might clam up in order to convince Alice that she can trust him. They both clam up. They both get small sentences. But then Bob knew all along that a big sentence was coming down the line, lethal injection, so he wanted to setup Alice to trust him enough so that Bob would have better odds at avoiding death. There are many cons that work by establishing "trust" with the mark by using small amounts of money, only to con the mark into handing over a much larger sum at which point the conman betrays the mark and flees. See the 1973 movie "The Sting" for an entertaining demonstration of building up trust using small amounts of money to extract an even bigger payoff later.

Games with limited iterations and iterations with different payoffs means that the players must be aware of their running totals of payoffs to determine if they have "enough" trust to trust the other player. If you, Bob, trusted Alice, and Alice returned that trust, and as a result, you Bob got an additional 20 million dollars in your bank account, then if the next iteration is for 50 million, you might want to note whether you can make up for the loss in later iterations if Alice betrays you.

The last few iterations of a long history of iterations may look like one-time games. Past behaviour is no guarantee of future results. The last iterations have no future moves to reward cooperation or punish betrayal.

The limited channels for communication and enforcement means that you can know only so much about the other player. Your enforcement mechanism is limited strictly to how you behave in future iterations, i.e. whether you stop cooperating with the other player or not.

In a population that is mostly fair, the population in general is better off with iterations than with one-time games. This reinforces the population to trust the rest of the population. But that can be a dogmatic strategy. The payoff matrix hasn't changed. The sucker's payoff remains.

Laissez Faire Capitalism Modeled as Multiple Iterated Prisoner's Dilemma

Laissez faire capitalists don't want any kind of government regulation on the market.They assert that the free market, the unregulated market, operates as an iterated prisoner's dilemma and will naturally tend towards the best possible solution for both players or all players.

Before game theory was developed and specifically before the iterated prisoner's dilemma was created, laissez faire capitalists said that this tendancy towards a natural solution was due to the "invisible hand" of capitalism. Since game theory has been developed, they will point to an iterated prisoner's dilemma and assert that prisoners will naturally tend to cooperate.

Alice wants groceries. So she goes to one store and buys some groceries there. Turns out a lot of the food is spoiled, or past due, or mislabled. Alice tries a different store. Again, she doesn't like the quality. Alice tries a third store and likes the quality of food and the prices. She starts shopping theire repeatedly. Years go by. The store changes hands. The new owners change things around. And Alice no longer likes the quality of food. She goes to a different store.

Alice tried a couple of stores who didn't "cooperate" with her. They tried to "rat" her out. So Alice kept going to different stores. When she found one she liked, she kept shopping there. Alice and the store went through many iterations of transactions. As long as the store delivered good food at a reasonable price, Alice continued shopping there. If the store stops, Alice moves on to a different store.

That's the stereotypical example of how the free market, buyers and sellers, maps onto the metaphor of the prisoner's dilemma.

The problem is that many transactions are really one-time interactions between people. You may buy groceries once a week, but how often do you buy a car? A house? A computer? How often do you get cancer and need to hire extensive hospital care? Those may occur more like one-time transactions. One time prisoner's dilemmas.

Another problem is that some transactions have such terrible negative consequences compared to the small price to purchase that it becomes difficult to manage the amount of trust. You may be able to buy a plane ticket to Las Vegas for only $25. But if the airline scrimps on maintenance and overworks their pilots to save money to be able to charge such a low price, you might end up dead, while the CEO gets to swim in a pool filled with your money.

Another problem is if you have iterated transactions with some company, but you don't have any real alternative to go to if you don't like the service. Two-year agreements with cell phone companies lock you in regardless of how bad their service gets during that time. The problems associated with changing your phone number used to lock customers to their carrier until the government forced companies to allow transfering phone numbers to other companies. You may be paying your cable bill every month, but many customers don't have any realistic alternatives.

Yet another problem is that you STILL don't know what the other player's motivations are. Are they offering you a good product at a good price because they're doing business that efficiently? Or did they have some cash saved up and are attempting a short price war with a competitor with no cash available in an attempt to run the competitor out of business? If you knew all the motivations, you might not purchase the under-priced product because you would want competition to remain strong. But the prisoner's in the dilemma must make decisions based on not knowing all the information.

So, the problem is that while in an infinitely iterated prisoner's dilemma between Alice and Bob, they will both find advantage in cooperating, the real world is a lot of one-time games and limited iteration games between many different, changing people. In the real world, the payoff isn't always "getting customer's money" but may be "getting rid of competition". In the real world, if you don't have alternatives, then you're a captive buyer, which means you can't go anywhere, which means you can't punish them for betraying you, which means they will treat each transaction as a one-time game, which means they stop cooperating with you the customer, because there is no long-term repercussions. In the real world, dying is a potential cost which is difficult to translate into a price the sick person is willing to pay to avoid.

Economic Inefficiencies in the Free Market

Economic efficiency is a measure of how much the production of goods and services can be increased with a fixed amount of available resources. If production cannot be increased, the economy is 100% efficient. If production can be increased using the same resources, the amount of increase reflects how inefficient the economy was.

Probably one of the biggest fallacies about the free market is that it is the 100% most efficient way for an economy to produce goods and services with some available set of resources. In fact, laissez-faire defenders often imply that any inefficiency in the current economy is due to government regulation. This implies that if you removed all government regulation, that the economy would be 100% efficient.

This doesn't add up in any mathematical sense.

Looking at market transactions as an iterated prisoner's dilemma, Alice must first first try some markets before finding a store that sells the product she wants, at the quality she wants, and a price she wants. Every store she spends money at that turns out to be not good quality at good price is money Alice wastes. If Bob decides to short-change Alice to meet his numbers for his quarterly report, then Alice wasted money.

If Alice dies because Bob insisted on not spending a penny more on safety, well, that's an inefficiency of massive cost.

An iterated prisoner's dilemma between the same two people over an infinite amount of time may tend towards the best possible outcome with no inefficiencies. But the real world consists of a lot of people having one-time and short-term-iterated transactions. Communicating with money is inefficient. Buying shoddy product before finding a good product is inefficient. And in the case of a one-time transaction, the seller may very well tend to "rat" on the customer, and neither one gets a good deal, and that too is inefficient.

The prisoner's dilemma naturally tends towards the second worst outcome available and the worst "fair" outcome available. In the prisoner's dilemma, a better outcome is available within the payoff matrix, but the problem is getting the players to choose that outcome without knowing what the other person will choose. You cannot accomplish this with a player's strategy, because strategy doesn't change the matrix.

Change the Payoff Matrix

The only way for two players in the prisoner's dilemma to choose the best, most efficient, outcome, when they don't know what each other's choice will be, is to change the payoff matrix.

In market transactions, changing the payoff matrix can be accomplished by government regulation.

Regulating the Prisoner's Dilemma

Alice and Bob and the rest of the population of the Land of Mervin are subject to a number of one-time prisoner's dilemmas. The difference is that the police pick up random people each time. It's not always Alice and Bob. One time it was Bob and Fred. The next time it was Charlie and Tom. The next time it was Helen and Kyle. But then the next time it was Bob and Fred again.

It isn't exactly an iterated game so cooperation can't be enforced via future turns. It results in inefficient outcomes.

But it it's something that the people realize is happening over and over and they want to deal with it. They want a more efficient outcome. Since it's occurring as one-time games, everyone is ratting out everyone else, and everyone is getting 5 years in prison. If everyone clammed up, then everyone would get only 1 year. But since it's not iterated, you can't use future turns to reward cooperation.

So, everyone in the Land of Mervin gets together and decides that they'll set up their own prison and legal system and create a law that says anyone who gets picked up by the cops must clam up or serve 30 years in Mervin's own new prison. The choices look like this:

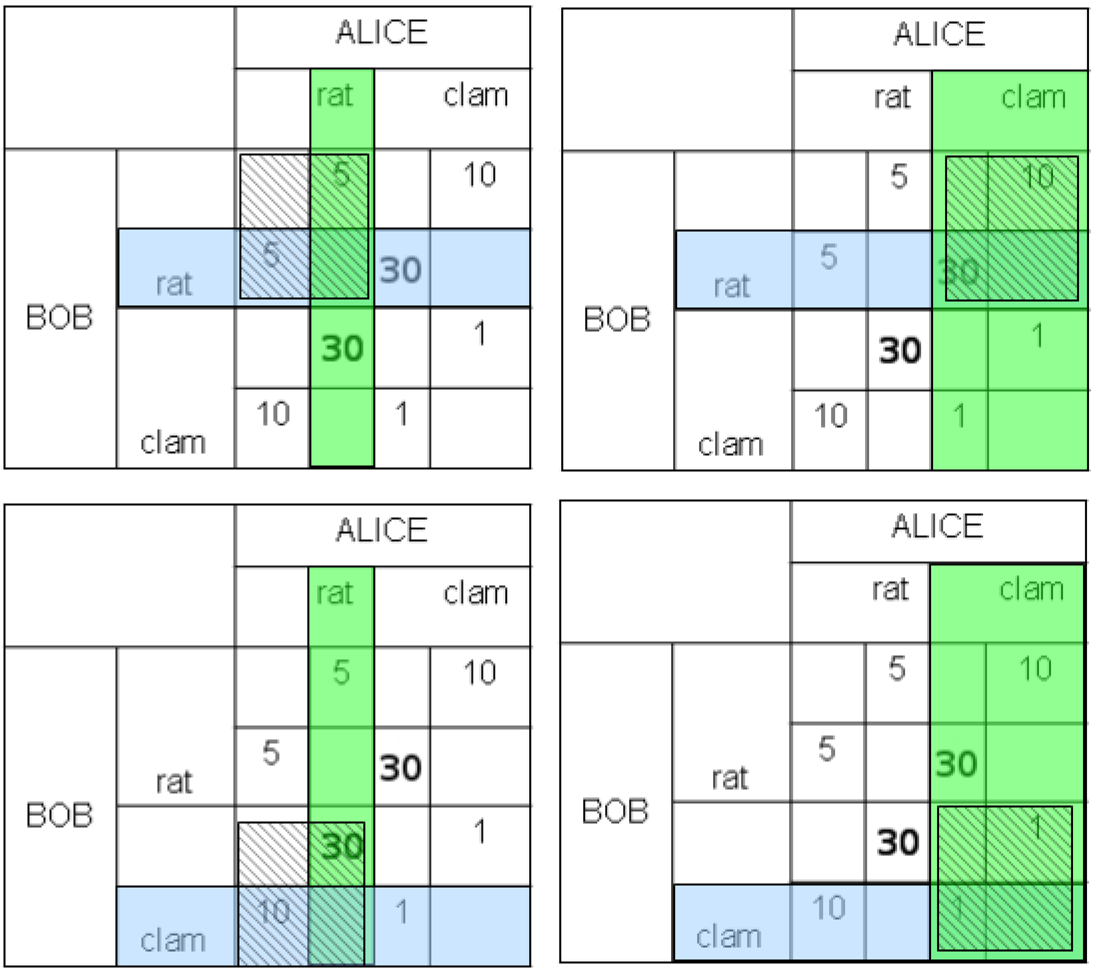

You, Bob the prisoner, still have 2 choices, clam up or rat.

Rat => 5 or 30 years in prison

Clam => 1 or 10 years in prison

If you want to minimize your time in prison, you'd choose to clam up. Whoever the other prisoner is would choose the same thing. And if you both clam up, you both get out in 1 year.

There are four possible sentences: 1, 5, 10, and 30. And regulation results in the shortest possible sentence for everyone. This means the regulated game is "stable". Both players operating for selfish reasons will choose to clam up to get themselves the shortest possible sentence.

The regulated prisoner's dilemma is just as efficient as the most efficient unregulated game, 1 year per player per iteration. But the regulated market has the advantage that it achieves the most efficient outcome even in a one-time game or a short-iterated game.

The unregulated one-time and short-term iterated game are inefficient because players will tend to not cooperate and get 5 year sentences each. That's an inefficiency of 4 years per player per turn. The regulated game is much MORE efficient than the unregulated game for one-time and short-term iterative games.

Overall, the regulated game is MORE efficient than an unregulated game with one-time, short-term iterated, and long-iterated games all taking place.

Now there will be a cost overhead associated with the Land of Mervin maintaining it's prison and legal system and so on. But that will be distributed to all the people of Mervin via a tax. Which means that it is a "fair" outcome too. Everyone is paying a few dollars a year to maintain the regulatory system, and everyone unfortunate enough to be picked up by the rogue police and subjected to the prisoner's dilemma will all cooperate and only serve 1 year.

In the unregulated version, everyone saves a few dollars in taxes. But people picked up by the rogue police are serving 5 years in prison. Because no post-release enforcement policy is allowed.

In either case, the same number of people are subject to the prisoner's dilemma, which means the regulated version is more efficient than the unregulated version.

Fallacy of Quasi Regulation in an Unregulated Market

The laissez faire defenders will often attempt to assert that their unregulated market actually has some form of quasi-regulation or another that is just as good as government regulation.

Reputation System

One example is the "reputation system". The laissez faire defenders will assert that the world would be just as good as it is now if we got rid of something like the FDA and allowed customers to "regulate" sellers via some kind of "reputation system". What they're really acknowledging is that many transactions are in fact one-time and short-term iterated transactions. Their "reputation system" attempts to convert the individual transactions into one infinitely iterated game. If Alice has a bad experience with Bob, and if Alice tells everyone about it, then Charlie and Dave and Eve would avoid Bob too. But in reality, that means you would find the best medicines by reading reports from the next of kin of customers who died from bad medicine. But snake oil would "naturally" be squeezed out of the market by this reputation system, so claim the laissez faire defenders.

The only problem I see with that is all the people who have to die before some medicine gets a bad enough "reputation" that it actually goes off the market. Rhino horn is still worth a lot of money to some people. Personally, I'd rather have everyone pay a few dollars in taxes and have an FDA organization instead of implementing some quasi drug regulations via a darwin-award kind of system crossed with youtube videos of the deceased.

Regulating bad medicine by having the customers die off is moronic, dogmatic, and even inefficient.

Boycott

Another quasi regulation system is "stop purchasing products from cheating companies". It doesn't work. If the last transaction between Alice and Bob is that Alice paid Bob twenty-thousand dollars and Bob only delivered one-thousand dollars worth of goods, how can Alice get that money back by buying stuff from Charlie? She can't. And her only laissez-faire recourse allowed is to start a boycott of Bob trying to convince everyone not to buy from Bob.

No Lawsuits ("Tort Reform")

Laissez faire defenders are so enamoured with how much power they believe exists in this wet noodle of a weapon called boycotts, that they often argue that to stop buying product somewhere is so powerful that no other post-transaction enforcement/corrective system is needed. They will often argue for "tort reform" which is a buzzword often used by laissez faire defenders to say "if they screw, you can't sue."

Lawsuits are an actual, existing, working post-transaction enforcement policy. And many laissez faire defenders want to get rid of it and replace it with reputation systems and boycotts. I assume this is because they are approaching the prisoner's dilemma of the economy with a dogmatic strategy of "always cooperate". Well, that's fine for Fred the Free Market evangilist and Larry the Laissez Faire fan. But not everyone follows that dogma. And if Alice gets shortchanged by Charlie the conman, then it's a bit of a problem if Fred and Larry tell Alice that she can't sue Charlie.

Invisible Hand

The "invisible hand" is another quasi regulation concept. But how it works is waved away with a lot of smoke and mirrors. Mathematically, non-cooperation is the default answer in one-time and short-iteration prisoner dilemma games. This is highly inefficient. But it is also mathematically proven to be the choice the player would make given that they have limited information about how the other person will behave and no real means of post-transaction enforcement.

Anyone invoking the invisible hand cannot do simple math. They're doing the equivalent of insisting that one plus one equals seven, and cooperation magically appears where it doesn't exist.

Unregulated Health Care as a One Time Prisoner's Dilemma

Unregulated health care often operates as a one-time prisoner's dilemma.

From an insured patient's point of view, you may be paying for medical insurance over a long period of time, but you don't actually see whether the insurance company will cooperate with you or rat you out until you get sick. When you go into the hospital for cancer treatment after paying insurance for twenty years, that's when you find out whether the insurance company will hold up their end of the deal or throw you under the bus.

If the insurance company decides not to pay for your treatment, your only recourse may be to attempt to sue (but tort reform would stop people from being able to sue). But sueing will cost you money and you don't know if you'll win. And you may be dead from cancer. Or you may be bankrupt because you're paying for your own cancer treatment.

Which means you spent twenty years cooperating with the insurance company, and when it came time for them to choose their action, cooperate or betray, they betrayed you. This is what is to be expected of a one-time prisoner's dilemma game. The players tend to rat each other out. And like a one-time prisoner's dilemma, you have no post-transaction enforcement.

A 2009 study of some of the larger insurance companies in California found that 20 to 40 percent of all medical claims were denied. Of all the bankruptcies filed in the United States, half of them are medical bankruptcies. People with massive medical bills they couldn't pay off. Of all those bankruptcies, 80% were filed by people who HAD MEDICAL INSURANCE but the company didn't pay for them.

From the point of view oft the uninsured, the circumstances create problems there too. If a person cannot afford insurance, the only option is to not get insurance and hope they don't get sick. If they do get sick, they usually have to wait until they are really sick. Sick enough to go to the emergency room. Then the hospital sends the person a huge bill, because emergency room treatments are a lot more expensive than regular health care. The person was too poor to pay for insurance and is definitely too poor to pay thirty grand for an emergency room visit. The person can't pay the bill. And the hospital passes the costs along to all the other patients.

In the case of the insurance company not paying the bill, the insurance company is doing solely what's in its own best interest and betraying the customer. The patient gets a huge bill, they can't pay it. They declare bankruptcy, and the hospital has to pass along the unpaid bills to everyone. In the case of the person who can't afford insurance, they're unable to cooperate because they can't afford insurance, and if they get sick, they can't afford to pay the bill, and the hospital is forced to pass along the costs to everyone else.

Either way, the fact is that there is no health care regulation that mandates cooperation, and its the hospitals who end up with the unpaid bills, which they then have to pass along to the rest of their customers. The lack of cooperation affects everyone who gets sick. And as a whole, the health industry operating this way is economically inefficient on a massive scale.

To improve efficiency, the approach is the same as any iterated prisoner's dilemma: regulation.

Regulating Health Care to Get the Best Fair Solution

Regulating any prisoner's dilemma to change the payoff matrix and make the best outcome the stable outcome improves the efficiency of the entire game. How would this be done with health care?

We need regulations on the insurance companies to prohibit them from betraying customers who make claims. Preexisting conditions should cease to exist as an insurance topic. Companies should not be allowed to reject you due to any preexisting condition and should not be able to charge you more for any preexisting condition.

This allows people to transfer from one insurance company to another more easily. Which is the direct way for customers to react to the quality of service they get from their insurance companies. It's post-transaction enforcement. Treat a customer poorly and they should be able to find a different insurer. And if a customer does get thrown under the bus by an insurance company and refuses to pay their bill, the customer shouldn't have that illness suddenly become a preexisting condition that makes it impossible to get insurance anywhere else. This is like government requiring phone companies to allow customers to transfer phone numbers. It encourages people to move from bad carriers. It rewards good business with more customers and punishes bad business with fewer customers.

Insurance claims should be paid within a certain amount of time. A company that drags their feet to avoid paying a claim should face the possibility of stiff fines if they don't hold up their end of the bargain. A company that recieves a claim and then tries to jetison you as a customer because you failed to dot some I or cross some T on a form three years ago, should face large fines. What's been happening is that the customer has been paying insurance for years, and when they make a large claim, the company goes through the records and finds an excuse to jettison them. They don't really care about the I's being dotted, it's just a retroactive excuse they invented to avoid paying the claim.

The fine should be large enough that it's cheaper for the company to pay the medical claim than to pay the fine.

Insurance companies have invented all manner of excuses to jetison customers. All these practices should be outlawed. There should be no lifetime limits and no annual limits. No preexisting conditions. And paperwork snafus should not be a justification to drop the customer. That would be a good start on the regulations regarding the insurance companies.

As for the uninsured, the young adults who dont' think they need it, the small business owner who can't afford it, and the unemployed person who is trying to pay for food, a public option would put them under some form of coverage, which would prevent health issues from being ignored until a visit to the emergency room and massive bills that no one can afford get racked up, and passed onto the other patients in the hospital.

Regulate insurance companies to reward their cooperation and punish their betrayals. Use a public option to insure those who can't afford it otherwise. Together, these two approaches would fix some of the larger problems plaguing health care.

Health care reform would cause both individuals (customers and patients) as well as insurance companies to cooperate, achieving a much more efficient result than the unregulated market is seeing now.

It's All About Trust

From a psychological point of view, it seems that a person's level of trust correlates to whether they tend towards cooperation or self interest in a prisoner's dilemma situation. The amount of trust also tends to correlate to where they land on the political spectrum. An absolute laissez-faire defender will tend to indicate signs of paranoia to some degree. They view any government as an affront to their liberty. A moderate-conservative will tend to not trust others and will tend to be worried about others taking advantage of them if they cooperate. They are often concerned how hypothetical welfare queens will take advantage of their cooperation and use it to buy a cadillac. A moderate-liberal will tend to trust others and will tend towards cooperative solutions. They also tend to support government regulation to prevent selfish players from taking advantage of their cooperation.

When approaching a real-world scenario that maps to some sort of prisoner's dilemma game, especially one-time or short-iteration games, the best solution is to implement regulation to keep players cooperating and everyone achieves the best fair outcome. But politically, people will oppose this for various reasons. They may present all manner of arguments that sound as if they're based in economic costs and benefits, but often their opposition really comes down to their lack of trust. They don't trust the players to cooperate. They don't trust regulations to work. They don't trust the government not to enslave them. They don't trust. The whole "Deathpanel!" propaganda around health care reform does not reflect any factual reality about regulation, but instead reflects the person's complete lack of trust in mankind. They want control to be in their hands. They want to make their decisions. They do not trust cooperation and want to act in their own self interest even if it means they achieve the worst possible outcome.

The driving force behind opposition to the best fair outcome is distrust. If someone wants to implement regulation for a one-time prisoner's dilemma situation, but is faced with opposition from a group of people, disproving their immediate arguments may not be a sufficient solution. Their immediate arguments may be nothing more than an outcome of the fear they are feeling due to their inherent distrust of others. If you want to achieve cooperation, especially regulatory cooperation which is the most efficient outcome, then you may have to make a concerted effort to reassure them to dispel their emotional fear and to encourage them to trust.

Summary

The prisoner's dilemma puts into quantifiable terms the issues that occur when two people are attempting a transaction but they don't know how the other person will act. Not knowing how the other player will act means that the player must choose to act in his own self interest. However when both players do this, it results in the worst possible outcome.

Cooperation can achieve the best possible fair outcome for both players, but there is no way to guarantee cooperation with a one-time prisoner's dilemma game. Iterated prisoner's dilemma games can use future games to reward cooperation and punish selfish players. But much of the real world involves one-time transactions and few-iterated transactions. And when cooperation is assumed but not enforced, conmen can operate at a large profit to take advantage of cooperative players.

Regulation can change the payoff matrix so that the lack of information about how the other person will act becomes less relevant. With regulation in effect, you know that if the other player does not cooperate, that player will end up paying a regulatory penalty. You know that the other player has incentive to cooperate. You know you have incentive to cooperate. Therefore regulation tends to result in the best fair outcome for both players, and regulation achieves this in one-time, short-iterated, and long-iterated games. An unregulated prisoner's dilemma can only expectc to see cooperation during a long-iterated game.

The efficiency of an unregulated game is a reflection of the fact that an unregulated game tends towards non-cooperation which results in the worst fair outcome. An unregulated, one-time prisoner's dilemma will tend to result in both prisoners not cooperating and both getting 5 year sentences. A regulated one-time prisoner's dilemma will tend to result in both prisoners cooperating and both getting 1 year sentences. This means an unregulated prisoner's dilemma game is unefficient by 4 years per prisoner per turn compared to the regulated game.

A person's trust regarding a real-world prisoner's dilemma scenario can be indicative of where they land on the political spectrum. Extreme right individuals don't trust anyone, tend to not cooperate, and view government as an afront to their liberty. Moderate right individuals don't trust that someone might take advantage of their cooperation if they were to offer it. Moderate left individuals tend to trust, and tend towards cooperative solutions. They also tend towards government regulation to achieve the most efficient and most fair outcome.

Regulating one-time and short-iterated prisoner dilemma games achieves the most efficient and most fair outcome. But to achieve regulation, you must have political support. But people who tend to distrust others will tend to oppose regulation and tend to support self-interest, non-cooperative solutions instead. They may present various arguments to disprove a regulated cooperative outcome. And while it is important to disprove the errors in these arguments, it may not be enough. You may have to find a way to reassure them, dispel their fear, and win their trust.

Or at least enough of them to get regulation passed. The pure laissez-faire paranoids are probably beyond saving.