Human Games

by C. S. Schroeder

The well established results of game theory indicate that in cases involving two players, and interests otherwise opposed, being rational one should decide upon the strategy of outing your complicit adversary in the following game: you are a representative of your nation at a tribunal; your adversary is a representative of another nation; you stand before a mediator in the on-going dispute between your countries. You each have a secret: that you have been running tests of weapons of mass destruction, against international laws; and you each know that the other knows your secret (and knows that they know, that they know, etc.). The best scenario, for you, in your estimation and the estimation of your adversary, is that you "out" your adversary, but your adversary does not "out" you. The worst scenario is that your adversary "outs" you, but you do not "out" your adversary; meanwhile, if you both choose the same, then you are better off if you both keep quiet about the WMD's. The "payoffs" can be represented in the following diagram:

B\A	Quiet	Out
Quiet	2\2	0\3
Out	3\0	1\1

Looking at this, you may feel that you, as a representative of A, should collude with the representative of B, to keep your mouths shut. But notice that whatever you agree to do, you are better off choosing to "out" your opponent, since if your opponent remains

quiet and you out your opponent, you gain 3, as opposed to 2, and if they out you, you gain 1 as opposed to 0, if you out them as well. In this model, it is a matter of mathematical fact that you are better off outing your adversary. But this model does not appreciate important subtleties in the situation. For this model to be accurate, you need to ignore the fact that in any sort of typical courtroom setting, there is an element of reaction. For instance, if you out B, then you can well expect the representative of B to react by outing A, since by doing so – in terms of the models payouts - they stand to recoup 1 point; moreover, if you are the first to do the "outing", you face the problem of risking your reputation as one who can't keep a secret and doesn't keep a promise. Factoring in these elements, the payouts may look as follows:

B\A	Quiet	Out
Quiet	2\2	1\.5
Out	.5\1	1\1

But realizing this, you would do better in your "outing" to out you both, as a measure of being "forthright" in your outing; and the payouts may in turn fall to something as follows:

B\A	Quiet	Out
Quiet	2\2	1\1
Out	1\1	1\1

This is assuming that what is gained in being forthright is lost in breaking the pact. But whether your adversary claims to experiment or not, the truth is out, and they will be caught, leaving them with the same payout of 1. In this better model of a more

typical case, it is clearly in the representatives best interest not to out the other party. They stand to gain more if they don't and no games-men-ship is going to help them if they do.

****

Now enter the island nation of Bali. While you and your adversary are world figures on the world stage, a third figure X shows up. X and his nation have been experimenting with WMD's as well, but frankly, they have no use for them; as such, when X shows up, he's got little to gain by keeping his mouth shut, and quite a bit to gain, relatively speaking (attention for their domestic issues, etc.), by outing the other two – provided they both hold their tongues.

X\ A or B	Quiet	Out
Quiet	0\old game	0\old game
Out	3\old game	0\old game

Surely they have other battles to fight? So it can likely be expected that they would concede and give the third player something to lose, right? This strategy works in the short term for you and your adversary from B, but it only invites everyone and anyone to step into the fray and claim their prize – and in the case of WMDs, that only means proliferation. Can the world, it's games, and it's players be so dysfunctional that extortion is the rule rather than the exception? Of course, morals does not math make, but it stands to wonder if through the artifice of game theory we have not deluded an ordinary man into feeling rational in his deprivation. What is missing, first and

foremost, is a deeper understanding of games and their payoffs.

****

By and large, the simple grid model of games is a sufficient representation of one-off betting games, but it fails to appreciate the subtleties of real games. What I have naively wondered my whole life is this: is there even such a thing as a good for you that is not also good for me? It is a very simple question that has even found its way into sitcom television, but on the face of it, the answer is strictly, yes. In such-n-such game (business, politics, chess) I win and you lose, given such-n-such choices. But no matter how well defined the moves and payoffs appear, there is always the fear that this is the battle won, which loses you the war.

The fear of unintended consequences is a very legitimate one. The nuclear arms race is a case in point. The U. S. during WWII was in a race to get the nuclear bomb first; all of the brightest minds available were collected in Los Alamos to put it together. They won the race – though it was later determined there was no race at all – and went on to beat Japan into submission with their new artillery, helping to put an end to the war. Meanwhile, a pair of relatively insignificant physicists (*qua* physicist) passed the blueprint on to the Russians, and in the end we had fifty years of nuclear fear. If there is a principle here, it is this: some things shouldn't even be figured out.

Too often we are zealous to prove our intelligence to the world. This makes bright young minds eager to play in grand games. In these games, their missions are handed to them, and their task is before them: solve this problem, advance in reputation, and repeat until you retire with acclaim. They operated most importantly as pieces in the larger games; but day to day Oppie and Johnny solved their puzzles for

their rewards, failing to acknowledge the larger games they were in. It was only when their contributions were over and fame assured, that any seeds of doubt grew into trees. (At least for Oppie; Von Neumann was without remorse, and given his grand pet project of controlling the weather, probably assumed we could engineer our way out of any mess later – perhaps he underestimated the rarity of his own talents).

Some will see this as a call for greater secrecy, but we must be aware of the illusions of the protections of privacy. On top of outright theft, of course, is the remarkable fact of symmetry in problems solved. It happens time and again that the same point of history yields similar solutions and theories overcoming the same "obstacles". This is frightening if only due to the fact that even should one figure something out, burn the paper, and drink a whole bottle of scotch – people everywhere may be figuring the same thing. There is no fearing a 'leak' in that case; the fear is that others will pursue it and use it against you, together with the feeling that you have no control over that, but to pursue it yourself. It has always been assumed that not pursuing the feared idea was a leap of faith. What I want to show here is the kind of game you may be in, such that it is rational not to pursue the feared idea.

****

Suppose for a second that you are Oppie. Oppie has an adversary, who in their personal battle would like to see him lose; in fact, Oppie's adversary has stuck his reputation on the idea that Oppie will fail if and only if he slanders him. Oppie's issue is always keeping his focus. While Oppie's adversary can only really feel pride in Oppie's defeat if he can believe that his slander caused Oppie's lack of focus:

Adversary\Oppie	Focus	don't
slander	-1\1	1\-1
don't	0\1	0\0

Oppie knows, then, that he will 'succeed' as long as he is focused. Of course, Oppie is part of a larger game, which he is aware. That larger game is the race to get the nuclear bomb for the US. As Oppie sees it, the science is nearly all in place, and it is a matter of engineering now, so if the team performs efficiently, it will get the bomb first. Oppie's adversary in the game is Heisenberg. Heisenberg is smart, but behind. Heisenberg needs to perform remarkably well to get the bomb first. So the payoffs seem as follows:

TeamHeisenberg\ Team Oppie	Efficient+	Efficient-
Extraordinary+	-10\10	10\-10
Extraordinary-	-10\10	0\0

Where the points represent utility for the development team.

Now, to the adversaries doing battle directly; the payouts are as follows for the actual demonstration that they have the bomb (assuming that if they both get there eventually, the U.S. will get there first and reap the reward for demonstration.)

Team Empire\ Team Allies	drop	don't
drop	-100\100	100\-100
don't	-100\100	0\0

To Oppie, it is all very clear. Game A is embedded in Game B and Game B embedded in Game C. Everyone Oppie identifies with stands to gain if he only remains focused.

It is arguable whether these payouts are accurate, but they are close enough for our purposes. The better question in this is why does Oppie identify with this team and this game?

Of course, I do not ask this question implying that Oppie should have defected to the other side. My question is, why would he identify at this level of resolution. Among games between people, there has hardly been a bigger stage than that depicted above; but Oppie is a part of still bigger games. Oppie's actions have consequences in these games as well – but he is hardly aware of these games at all.

Nature\Humans	Nucs	No Nucs
cataclysm	-1000\-1000	-1000\-1000
none	-1000\1000	0\0

Or,

Universe\Life	Human Nucs	No Human Nucs
cataclysm	0\-100000	0\-100000
none	0\0	0\0

Even keeping in mind that the actual numerical payouts are conjecture, certain folks will roll their eyes at this point. Nonetheless, these remain valid games: certainly, humans have fears of the living natural world outside of themselves and take measures to protect themselves against it; in modern culture, these are fears of viruses, bacteria, and vectors for genetic disease. And life in general has their own fears of the non-living, chaotic natural world. The issue can all come to light in a doomsday scenario as follows: suppose that nuclear radiation from artillery led to genetic mutations in

bacteria, which led to widespread and uncontrollable disease? In this case, life would lose out, humans in turn lose out, and the games below these in the hierarchy would have trivial payouts by comparison.

****

The biggest obstacle to accepting the relevance of such games to decision making is the sense that we cannot know their payouts. The reasoning is: without knowledge of the payouts, we cannot possibly know the correct ways to play these games; if we cannot know the correct ways to play these games then we do not have control over their outcomes; we cannot concern ourselves with what is outside of our control, so these games are irrelevant to decision making; we must therefore focus strictly on what we know, that if we win the race, our teams will be better off. I want to argue here that this kind of thinking, which I believe to be at the root of much trouble, is not inevitable.

If the issue is what Oppie really should have done, things are rather complicated. In game theory, if your opponent is not rational, then all bets are off, since they may decide to abide by the payoffs or not, at their fancy. Given Hitler's undoubted instabilities, one can only say that the opponents rationality must have been suspect, and not assumed. On the other hand, we assume outright that Bali knows what they are doing, as a nation of intelligent people, who see this opportunity as a chance to make things right. And we assume that the intelligent paranoids with their dream-fears are, despite their paranoia, rational for all that. So whether the terrible idea should come from hope or fear, the question remains, why should it not be pursued?

While the game these players see themselves in are apparently of foremost importance, they must be *mindful* of the other games they are in. I do not use this term

"mindful" in a flippant way. In fact, I believe that mindfulness just is being aware to your greatest ability of the higher consequences of your actions *when you do not even know what games you are in.* The greatest obstacle to mindfulness is respecting the higher games you are a part of when you do not know what those games are about. How can you be aware of the consequences of your actions when you cannot analyze the higher games?! What needs to be appreciated are the connections between games. What I've hoped to bring to light is that the games you play are tied to the games your higher organizations play as well – forming a hierarchy of games (or at least a lattice of some form). Ascending the hierarchy there are innumerable games which you are a part, and to which your actions contribute, but which you do not even approximately understand. Individuals can expand their window of association to discover how greedy and risky behavior effects the organizations which they are a part, through proper methods of concentration meditation; but what we ultimately seek is a general analytical framework under which individuals and theoreticians alike may achieve a deeper understanding of the relationships between games. We will but skim the surface here, but I hope give it a start here.

****

There are two vectors by which one game may influence another. The first is that of the influence that game play has on the strategies employed by participants in other, future games played by those agents, among themselves and others. The second is that of the influence of actions and outcomes on the structure of other games, those played by the same participants or otherwise. The first type of influence has been well studied and is perhaps the most important conceptually, but in reality it ignores a great deal of

significance. Concentrating solely on the former to the neglect of the latter is akin to concerning yourself with computability at the neglect of complexity; for should one be able to specify all of their games into one grand game, then perhaps the classical study of strategy would suffice, but this task is far too complex in space and time to achieve, so we must concern ourselves to do our best local modeling of the games we play, and calculate back in the potential external influence of actions and outcomes to a given approximation. The hope is to account for the statistical influence that our choice has on the structure of our other games; and how that change of structure should effect our current choice. This area of inquiry could be called *The Theory of Full Consequence.*

Typical examples that come first to mind are unintended consequences that are detrimental to some other cause of the agent. An example of this may be false incarceration; the unintended consequence of criminal proceedings can be rebellion on the part of the falsely accused and/or payouts to them once they are freed; note that the rebellion or payouts are here the relevant unintended consequences, for these are often not factored into decision making, even should some rate of false incarceration be known and assumed. Another may be industrialization itself, where the general effects of manufacturing on the environment are either ignored or assumed to be reversible, when one undergoes an enterprise; only to find that byproducts take decades to clean up. In these cases the actions by an agent in one game contribute to action(s) by an organization in another game; that organization is either an organization which the agent is a part of or an organization which is opposed by an organization the agent is a part of; and the actions that it contributes to are in either case detrimental to the purposes of the organization in that game, which the agent is a part of.

Of course, we cannot assume at the outset that all such consequences are bad. Should the consequences be arbitrary, we would expect them to sometimes be good,

sometimes bad, and sometimes neutral. As such, we call the study of the effects of such unintended consequences and their import to decision making The Theory of Full Consequence and not *Blowback Theory*. In general, the theory will not assume that such unintended consequences are bad; rather it will only assume that the unintended consequences are arbitrary. But what it shows is that:

Even assuming arbitrary unintended consequences, your decision making process should change, because your model of games is essentially limited.

****

Suppose that your organization specifies a game that it wants to play; as part of this game, you are assigned a 'smaller' game to play, in order to help the organization succeed in its larger game; your actions in this game are specified to some, incomplete extent, so as you play this game, your actions interfere with outcomes of other games and actions of other agents, who are outside of your engagement, in arbitrary ways. Assuming, simply, that the space in which we have to carry out our actions is constrained and crowded to some degree, then we can assume that there is some mean percentage of other games which your game will interfere with; I will call this mean percentage of interference *N*; and N multiplied by the fraction of these games which we, or our organizations are a party to, I call *B*. Interference with a game can occur in apparently three different ways.

First, it can directly modify the state of nature or psychologies in a way which changes the payouts of the game directly; for instance, if artillery ruin crops which were payouts in economic games; or economic moves to 'put money in people's pockets',

lead to inflation, which lowers the value of the dollar itself. The second way is that it can change the state of nature or psychologies in a way which changes what the outcomes - or the likelihood of outcomes - of actions are; for instance, if a CEO finds a sales pitch from a company's representative to be offensive, then the outcome of the sales pitch to be given by a different representative, on a different product from the same company is going to be negatively influenced. The third form of interference is that in which a game-play influences the actions in the game by either executing them, or bringing them nearer or farther from execution; for instance, if a sales rep for your organization lands a big deal with a company, that company could immediately decide to sign your deal, with the relationship between the companies now established, effectively completing the execution of an action on your part; on the other hand, if someone gives their password in plain text to a colleague in order to expedite a transaction in a particular deal, but in doing so, compromises information on other deals, then the act may be contributing to an adversaries efforts to collect information on your company; and finally, often resources are allocated to one game at the expense of another, moving one further away from executing good actions in the other game; such cases occur whenever there is 'an urgent matter that needs our immediate attention', and resources are subsequently diverted.

The general approach that I suggest here is to model these forms of influence as follows: we can effectively understand that the influence of one game on the actions of another as the addition or removal of steps to completing an action of varying cost, with the understanding that should complexity (“the number of steps”) go to zero, then the action is effectively executed. The number of steps that may be added or removed may be understood as having some distribution, and an arbitrary step may be understood as having a severity distribution. Second, we assume that the transition from action to

outcome (given action of the opposition) is influenced by other games as well, but also amounts to the addition and removal of steps of variable cost, so can be modeled similar to the effect on actions. Third, influences on your opponents actions and transitions from action to outcome (outcomes given your outcomes) can also be modeled as additions and removal of steps of varying costs, so that we have a single effective distribution for these vectors of influence between any two games. Finally, influences on the utility of outcomes can be directly modeled, as having some other distribution. And this completes the modeling of influence.

These distributions are to be used with an approximation of N and B for each choice in the game under evaluation, to see how the potential payouts from actions are to be modified in certain ways. Furthermore, we can work backwards to assess the cost to our game of interference from other games. All with the hope that in the end we can better understand the full consequences of our course of action, and of playing this game in the first place. This may sound like a daunting task, but general rules of thumb should be available for evaluating N and B, based on the "footprint" of the actions in the space of our other games.

****

The execution of an arbitrary action in an important game is rather frightening. If the remaining steps to execution are both cheap and not complex, then there is a risk of the action being executed in a game, as a result of other human games, independent of the actions themselves being decided upon in the game in question. This is the essence of risk among human games in our model. On the other hand, should an action be simple but not cheap in human terms, it may still represent a risk, only a risk accountable not

to people (in most cases) but things – e.g. earthquakes. In general, simple actions with low cost, but great influence, are the primary risk to ones games. It stands to reason that sufficient complexity is actually important for the mitigation of risk. Should we not want to give up our games.

In general, if one can evaluate N, B, complexity, and the cost/utility of action and outcomes, one will have a baseline against which one can evaluate the risk of actions in your game on other games and the risk of actions on your game from other games. This does not mean one should toss to the side matters of evaluating particular risks on your game and particular risks from your game. Only that one can go some way to evaluating risks analytically and independent of the specific risks that one may find in a dreamy paranoid search. The hope is that with such general methods a Theory of Full Consequence can be achieved.

I do not know what the end result of a full analysis will be, but I hypothesize the following. In general, it is the failing of humans to comprehend that when they decide upon actions in arbitrary small games, they are bringing their higher organizations, which they are always acting as agents for whether they understand it or not, closer or farther away to executing actions in grander games. Should they reduce the complexity of executing these actions too close to zero, then the actions are rather likely to be executed whether we like it or not, as a matter of chance influences from other games. As agents of higher organizations, including life itself, we must do our best to be mindful and respectful in our decisions. I have tried here to make this clear, while also providing some inspiration for handling these matters analytically – so that the cynic too can have some hope.

****

If we return to Bali, the point of pressure on fellow nations cannot be, now what will

you do for us. Bali itself must respect what is rational for their higher organizations, and in doing so the pressure should be, stop the relatively small games which bring us closer to calamity. If Bali should take this position, they can expect that done well and in earnest, they may enter the fold; but they must at any rate understand that public extortion cannot happen. An understanding of a complete *Theory of Full Consequence* is intended to help such decisions be effectively made and avoid the risks to society and life itself, which no man really wants to accept.