DECISION THEORY PRIMER

*If you're in a hurry, you should at least take a look at part III of this handout.

Part I: Decisions under Risk

Suppose you draw once from an ordinary deck of 52 cards, and you want to decide whether to bet on drawing a red card or bet on drawing a face card. The chances of getting a red card are 50%, and the chances of getting a face card can be calculated to be 23%. That makes it sound like it's better to bet on a red card. Yet suppose that you will be paid significantly more if you place a winning bet on a face card. Then, it may be more rational to bet on the face card, even though it has less of a chance of winning. But how should we decide such a thing?

The decision on which bet to place can be represented in the following table:

		STATES			
		Win (State1)		Lose (State	2)
		+1M T		-1M ⊤	
	Bet Red Card				
ACTS			50%		50%
		+3M T		-1M ⊤	
	Bet Face Card				
			23%		77%

Each square of the table indicates what the relevant monetary outcome will be. Thus, the first square indicates that if you bet on a red card and you win, you will get 1 million tenge. Whereas the second square indicates that if you lose on that bet, you will lose 1 million tenge. And so on.

One way to decide how to bet, in this sort of situation, is to calculate the *expected payout* ("EP") of each bet. This is basically the amount of money you can expect to gain *on average* from one of the choices. Thus, if we calculate the EP for each of the two bets, then we can see which bet will be financially better on average.

The EP is calculated as follows (assuming that the states are mutually exclusive and exhaustive):

 $EP = [Payout_{State1} \times Prob(State1)] + [Payout_{State2} \times Prob(State2)]$

Applying this formula, we can determine that the EP of Betting Red is:

 $(1MT \times .5) + (-1MT \times .5) = 500K + -500KT = 0T$

Whereas, the EP of Betting Face Card is:

 $(3MT \times .23) + (-1MT \times .77) = 690KT + -770KT = -80T$

A further assumption here is that you should do what has the greatest EP. So under that assumption, it deductively follows that you should bet red.

Part II: Decisions under Ignorance

Notice that calculating the EP requires you to know the probabilities of the states. Unfortunately, very often we don't know those probabilities. Nonetheless, there are still some reasonable suggestions about how you should make your decision.

Suppose for example you are deciding on whether to study to become a lawyer or to become a philosophy professor. Now if you study to be a philosophy professor, you're not guaranteed to get a job—so there is potentially a significant financial loss if you make that choice. But similarly—if you study to be a lawyer, it is possible that you won't get a job as a lawyer (such things are never guaranteed). And law school might be a significantly greater cost. But you may not know the probabilities involved here. That is, you may not have any idea how likely it is to get a job as a philosophy professor or how likely it is to get a job as a lawyer. That means our decision table will look something like this:

		STATES			
		You get job	You don't get job		
		+1M T	-5M T		
	Study Law				
ACTS		[unknown]	[unknown]		
		+1M T	_1M T		
	Study Philosophy				
		[unknown]	[unknown]		

However, despite our ignorance of the probabilities, we can still try out one guideline for making the decision known as the *dominance rule*.

Dominance Rule: Do not choose dominated acts.

An act A dominates an act $B =_{df}$ the outcome of doing A is always at least as good, and sometimes better, than the outcome of doing B.

Applying this rule to our decision, you should study to become a philosophy professor. After all, if you examine the decision table, you'll see that the outcome of studying philosophy is at least as good, and sometimes better, than the outcome where you study law.

However, sometimes the dominance rule does not apply to a decision. Suppose you have three choices on what to study:

	You get job	You don't get job
	+1M T	_1M T
Study philosophy		
	[unknown]	[unknown]
	+3MT	-2MT
Study engineering		
	[unknown]	[unknown]
	+6M T	-5MT
Study medicine		
	[unknown]	[unknown]

Here, no act dominates the others. No act has an outcome that is always at least as good as the other two. But decision theorists offer other proposals about how to make a decision in such cases. Consider for example:

Maximin Rule: Choose the act where losing would penalize you the least. ("Play It Safe")

If we apply *this* rule, then you should be a philosophy professor. (What if two acts have the least penalties? Then, you might apply the "Lexical" Maximin Rule: Do what also has the "second least" penalty. If there's still a tie, do what has the "third least" penalty. And so on.)

Nonetheless, there are other rules you might consider when you are ignorant of the probabilities:

Principle of Insufficient Reason: In your ignorance, assume the states are equally probable and treat it like a Decision under Risk. ("Ignorance means Equiprobable")

"EP" of being a Phil Prof: $(1MT \times .5) + (-1MT \times .5) = 500K + -500KT = 0T$

"EP" of being an Engineer: $(3MT \times .5) + (-2MT \times .5) = 1.5MT + -1MT = 500KT$ "EP" of being an M.D.

(6MT x .5) + (-5MT x .5) = 3MT + -2.5MT = 500KT

Hence, you should be either an Engineer or an M.D. (Doesn't matter which.)

Optimism-Pessimism Rule: Represent your level of optimism using a number $0 \le n \le 1$ (a larger *n* means greater optimism). Then, use *n* like a probability. ("Weight Payouts by Optimism Level")

OP-Payout of an Act A: (Greatest possible payout for $A \ge n$) + (Least possible payout for $A \ge (1-n)$)

Example Suppose your level of optimism is .7. Then:

OP-Payout of being a Phil Prof: $(1M\overline{T} \times .7) + (-1M\overline{T} \times .3) = 700K + -300K\overline{T} = 400K\overline{T}$ OP-Payout of being an Engineer: $(3M\overline{T} \times .7) + (-2M\overline{T} \times .3) = 2.1M\overline{T} + -600K\overline{T} = 1.5M\overline{T}$ OP-Payout of being an M.D. $(6M\overline{T} \times .7) + (-5M\overline{T} \times .3) = 4.2M\overline{T} + -1.5M\overline{T} = 2.7M\overline{T}$ *Hence*, you should be an M.D.

Example 2: Suppose your level of optimism is .3. Then:

OP-Payout of being a Phil Prof: (1MT x .3) + (-1MT x .7) = 300K + -700KT = -400KT OP-Payout of being an Engineer: $(3MT \times .3) + (-2MT \times .7) = 900KT + -1.4MT = -500KT$

OP-Payout of being an M.D. $(6MT \times .3) + (-5MT \times .7) = 1.8MT + -4.5MT = -2.7MT$

Hence, you should be a philosophy professor.

Part III: Retrospective on Decision Theory

The decisions we've just reviewed are odd: They assume that money is the only thing that matters, that you know the gains and losses for every act/state pair, and that either you know the probabilities exactly or have no idea at all. Real-life decisions rarely fit any of these assumptions. Still, Decision Theory provides *some* clarity on what it means to make a rational decision.

What's more, Decision Theory indicates some crucial "checks" that can lead to more rational decisions. First, if you make a list of "pros and cons" under each option, you are already doing much better than most people. In Decision Theory, this happens when we:

1. Specify all candidate acts and identify the possible outcomes under each act.

But crucially, Decision Theory takes the process further. It now asks us to:

2. Consider the *probability* of each outcome under each act.

So if we have a list of pros and cons, it might be good to consider *how likely* each of the pros and cons are, under each option. These probabilities may greatly affect how we view the options. For a given act, is the positive outcome a sure thing? Or is it more of a pipe dream?

The aim is ultimately to know what is the best "bet," considering the possible outcomes and their probabilities. This happens in Decision Theory when we:

3. Compare the probability-outcome pairs for each act.

The comparison tells us what is most likely to yield a good outcome and/or divert a bad outcome.

In sum, identifying outcomes, considering probabilities, and comparing the results is beneficial to decision0making, regardless of whether you use Decision Theory. These steps discourage you from deciding too hastily, i.e., without adequately considering crucial aspects of the decision.

Bonus: Decision Theory also offers *some* guidance on your choices. It recommends asking:

If you know the probabilities, which act has the *greatest expected payout (EP)*? If you don't know the probabilities:

Does any act *dominate* the others? Which act has the *largest minimum outcome*? Which has the largest EP if your *optimism level is used like a probability*? Which has the largest EP if the states are *equiprobable*? I would repeat that the answers here should not dictate your choice. Again, that would assume (falsely!) that finances are the only thing that matter. Besides, we saw that different decision rules often recommend different choices. Still, asking such questions may offer different perspectives on your decision and thus may help illuminate which option is most rational.

Nonetheless, we should recognize a deeper assumption at work. Imagine that you could somehow "measure" the expected monetary *and* non-monetary outcomes of each act. Imagine that for each possible act, you can somehow get a number that represents *the expected balance of all positive and negative effects* of the act. Then, if you are rational, you clearly should choose the act with the highest "expected outcome" in that sense, right?

Immanuel Kant (1724-1804) disagrees. His view is that rationality requires consistency, and that some acts are *inconsistent with our values* even if they maximize the outcome. For example, consider that you could kill a homeless person with no family or friends—and then collect their heart, lungs, liver, and kidneys, to save the lives of five people in the hospital. You sacrifice one life to save five—that's a great outcome! You still probably wouldn't do it, and Kant thinks this is rational. If you killed a person just for their organs, that would be *inconsistent* with your values, and so you could not rationally choose this, even if the outcome would be better.

Regardless, even Kant would agree that it helps to *consider* the likely consequences of your actions, even if those don't *determine* what is rational. A decision which does not consider carefully the expected outcomes still looks hasty.