MAKING GOOD CHOICES: AN INTRODUCTION TO PRACTICAL REASONING

GLOSSARY (updated 3/31/12)

The following definitions are brief versions of the fuller definitions provided on the page(s) referred to. Words and phrases in italics have their own glossary entry.

Action-consequence pairs – possible courses-of-action plus their consequences that form an *agent's* menu of *options*. Example: if your goal is to visit your brother who lives 500 miles away, your options might include these three action-consequence pairs: (i) driving-arriving, (ii) hitchhiking-arriving, (iii) bus-arriving, providing these three courses of action are possible for you to do. (p. 13)

Additivity – in *multi-criteria decisions*, one of six conditions for *attributes*. Additivity is a principle of value independence: for any two attributes, the *value* assigned to each separately must sum to the value assigned to both collectively. A set of attributes must obey this principle in order for each attribute to be assigned a value or weight that will add up to the total value the agent places on the goal. Example: if you separately value your desk-top computer without electricity a small amount x, and you separately value the electricity needed to run your computer a small amount y, then these two values will fail the principle of additivity if you value your computer running on that electricity greater than the sum of each separate value (which you typically will). (p. 80)

Agent – anything (e.g., person, institution, nation) making a *decision*; anything that has a *goal* and must choose from at least two *options* what to do to achieve it. Example: if a college sets a goal to build a new student-activities building and must decide which of five architectural designs to accept, the college is an agent. (p. 9)

Arbitration – a set of principles and methods for arriving at a fair *rational choice* within a *bargaining* problem. Example: the arbitration principle "share-and-share-alike" would assign equal pieces of birthday cake to three children each demanding a bigger piece than the other two get. (p. 258)

Asymmetrical game – any *game* in which it would matter to the *agents* which *option-outcome* position each would be assigned; given the options, the level of *goal* achievement is unequal according to which position an agent is assigned in the game. To be contrasted with *symmetrical games*. Example: suppose two agents are flipping pennies with heads winning a nickel and tails losing nickel for one agent (say, the one who goes first), and heads winning a dime and tails losing a nickel for the other agent (the one who goes second). These agents would not be (and should not be, assuming winning money is the goal) indifferent as to who goes first and who goes second. (pp. 177 and 246)

Attribute – a property or category that connects each *option's outcome* with a complex *goal's objectives*. For *single-criterion decisions*, the goal's single objective serves as the attribute, but for *multi-criteria decisions*, a set of attributes is required, each having a *value* or weight. Example: suppose your goal is to buy a good car, and one objective is a car using as little fuel as possible. You drive a certain car, say 100 miles, and an outcome is that it uses a certain amount of fuel, say ¼ tank of fuel. The attribute "amount of fuel consumption per distance" is the category linking this specific outcome (100 miles in ¼ tank) to that objective (minimum fuel use), and this attribute allows you to compare the outcomes (the specific amounts of fuel used) for each car you might consider buying. (p. 76)

Availability error – the mistaken *belief* in assigning initial *pure* or *factual probabilities* to a *state* that the more mentally available an event is to an *agent* (for example, easy to remember) the more probable it must be. Example: someone who was in lower Manhattan on the morning of

9/11/2001 might tend to assign an unrealistically high probability to the danger of being in a terrorist attack, based on the easily remembered trauma of that event. (p. 111)

Backward induction problem – a problem for cooperation. In repeated *potentially cooperative games* in which *agents* are considering to *cooperate*, but in which each could have better *goal* achievement by singly *defecting* or worse goal achievement by being *suckered*, it is rational for each agent to plan to defect in the last decision in the series. But each agent will want to "beat the other to the punch" and so it will become rational to defect in the next to last decision in the series. The same reasoning applied to each earlier decision makes it *irrational* to choose to cooperate in the first place. The benefits of cooperation will be lost, but only if the last decision in the series is known beforehand. Example: two people, A and B, agree to cooperate: A will lend B money at the beginning of each week, and B will pay A back at the end of each week. But A will not lend if not paid back, and B would like not to have to pay A back unless it's to be lent again. Suppose B learns that in 3 weeks A will die mid-week, and so plans not to pay A back at the end of week 3 (for there will be no lending in week four). A knows this and so plans not to lend B the money at the beginning of week 3. B expects this and so B plans not to pay A back at the end of this week 2. A realizes this and so plans not to lend B money at the beginning of week 2. B assumes as much, and so plans not to pay A back at the end of week 1. A expects as much, and so plans not to lend B money at the beginning of week 1. And so, by backward induction, A and B void their agreement and don't cooperate at all. (p. 228)

Bargaining – *practical reasoning* within a bargaining problem. A bargaining problem is a decision problem of dividing a goal among agents that gives each a fair share. Bargaining involves applying principles and methods of an *arbitration* scheme. Example: two house mates who hosted a big party last night must decide in the morning how to share the task of cleaning up the huge mess, assuming their goal is to clean up after the party. (pp. 246 and 256)

Belief – a propositional attitude on the part of an *agent* of accepting a statement as true with a *degree of confidence* or strength. A reasonable belief is a belief whose degree of confidence matches the *probability* that the statement is true (that is: the probability that things are just the way the statement says they are). Example: it is reasonable to believe with a .7 degree of confidence that it will rain (at your location) tomorrow if there is actually a 70% chance of rain (at your location) tomorrow. (p. 105)

Binding agreement – in potentially *cooperative games*, any agreement between *agents* that is enforceable in the event that an agent is tempted to violate it or fails to abide by it. An important step in *negotiations* is getting agents to agree on an enforcement mechanism or power that will serve to bind them to their cooperative decisions in case the temptation to *defect* or exploit one another increases. Example: agents who sign a legal contract are forced to keep to the terms of their decisions to cooperate by the penalties associated with breach of contract. (pp. 203 and 252)

Binding the will – a *fallacy of practical reasoning* due to personality weakness in which an *agent* does not sufficiently commit to or bind herself to a *rational choice*, and as a result gives up a challenging decision that the agent should remain in. Example: an agent who decides to quit smoking for reasons of health, ends up smoking again (did not sufficiently bind his will) yet still has good health as his goal. (p. 36)

Bounded rationality – a theory of practical rationality that seeks to justify methods of *practical reasoning* and principles of *rational choice* that agents actually use within the realistic limits and bounds under which real human decisions are typically made. Example: a decision arrived at by *satisficing* methods is a rational choice, according to the theory of bounded rationality, even thought it might fail to be a rational choice by standards requiring maximizing methods. (p. 162)

Brinkmanship – a *potentially cooperative game*, similar to *chicken*, in which each *agent* prefers a *free ride* to *mutual defection*, mutual defection to *mutual cooperation*, and mutual cooperation to

being a *sucker*. The *rational choice* solution for brinksmanship is mutual defection, representing equal *goal* loss for the agents. Such a game should be avoided, if at all possible. (p. 213)

Certainty – a *decision problem* in which the *agent* has a *belief* with a reasonable *degree of confidence* of 1.0 (that is: is sure) that, for each option, the *state* required for the *option* to have its *outcome* exists (or will exist at the time an option is acted on, if chosen). An agent might be certain by *default* or by direct evidence. Certainty is one of three decision situations, the other two being decisions under (conditions of) *risk*, and under *ignorance*. Example: typically an agent who decides to continue watching TV for another 15 minutes is (practically) certain that the TV, her eyes, her brain, etc. will continue working properly so that the decision achieves its outcome of continued TV watching. But if the agent is very sleepy, or if there is a thunder storm taking place, the agent cannot reasonably be certain that her decision to continue watching TV will result in continued TV watching. (pp. 17, 69, and 102)

Chicken – a potentially *cooperative game* in which each *agent* prefers a *free ride* to *mutual cooperation*, mutual cooperation to being a *sucker*, and being a sucker to *mutual defection.* Chicken suffers from an *equilibrium selection problem*, and consequently has no standard *rational choice* solution. Agents in the game of chicken must resort to non-rational ways, typically forms of deception and bluff, of coordinating decisions. Chicken should not be confused with a related game: *brinksmanship*. (p. 208)

Clash of wills – a potentially *cooperative game* in which each *agent* prefers a *free ride* to being a *sucker* (giving in), being a *sucker* to *mutual defection*, and *mutual defection* to *mutual cooperation*. If a clash of wills is *asymmetrical* it is a weak clash of wills and has a *rational choice* solution, but if it is *symmetrical* there is an *equilibrium selection problem* with no standard rational choice solution; agents must turn to non-rational ways (e.g., flipping a coin) to coordinate their decisions so that one (and only one) agent cooperates. A clash of wills is one-time if it happens between agents once, or iterated if the agents are in a series of repeated clashes of will. (p. 200)

Common interest – two or more *agents* that have an interest in the same thing(s), and that have an interest in each having an interest in the same thing(s), and that have an interest in each having an interest in each having an interest in the same thing(s), … , and so on. Common interest is a central feature of groups whose members are *social agents*. Common interests are contrasted with people having "interests in common," the latter referring to any two or more *agents* that happen to be interested in the same things, but may not even be aware of each other and have no interest in each maintaining an interest in the same thing(s). Example: a couple trying to make their marriage work have a common interest, not just an interest in common. (p. 10)

Common knowledge – in *interdependent decisions* (*games*), the assumption that each *agent* knows the decision situation, and that each agent knows that each agent knows the decision situation, and that each agent knows that each agent knows that each agent knows the decision situation, … , and so on. *Common knowledge* is contrasted with having "knowledge in common," the latter referring to two people who happen to know the same things, but are not aware that they do. Example: take two people playing cards with each other, each holding an ace. They have this piece of knowledge in common: (at least) one of us holds an ace. But this is not *common knowledge*, for each might think that she is the only one holding an ace, and so would believe that she is the only one possessing this piece of knowledge. But if both look at each other's hand and see the aces, and both know that they have each done this, then the knowledge: (at least) one of us has an ace, becomes common knowledge. (p. 169)

Competitive game – any interdependent decision situation (*game*) in which *agents* have complete conflict of interests, so that any *goal* achievement of one agent means an equal amount of goal loss for the other agent(s). Also called "zero sum games" because the *outcomes* for any pair of *options* sums to zero. Example: two breakfast cafés located near each other are in a

competitive game for the same pool of customers who can eat only one breakfast per day; a customer gained by one café is necessarily a customer lost by the other on any given day. (p. 170)

Complex goal – any *goal* having at least two parts (*objectives*) all of which must be achieved in order for the whole goal to be achieved. Example: if your goal is to earn a college degree then you have a complex goal that can only be achieved by achieving: admission to a college, passing all required courses, earning all required credits, paying required costs and fees, declaring a major and minor, etc. (p. 12)

Conjunction rule – the rule for combining probabilities that says to find the probability of two independents events (a and b) happening, multiply the probabilities of each happening; this product is the probability of both happening. P(a and b) = P(a) x P(b). Used in decisions under *risk* for multi-stage *states* of the world. Example: to find the probability that you will get both a flat tire and (independently) a speeding ticket while driving from Boston to NYC at 15MPH over the speed limit, multiply the probability of your getting a flat tire while driving from Boston to NYC at 15MPH over the speed limit times the probability that you will get a speeding ticket while driving from Boston to NYC at 15MPH over the speed limit. (p. 115)

Convention – a self-sustaining solution, or avoidance, of a coordination problem (a potentially *cooperative game*) that pre-decides *options* for *agents* needing to coordinate their activities so that *cooperation* results. Example: when going up or down a stairway, the convention is to stay to the right, rather than having to deciding between passing on the left or the right each time a stairway is used. (p. 203)

Cooperation – in a potentially cooperative *game* in which an *agen*t has two *options*, cooperate with the other agent or don't cooperate with the other agent, cooperation is choosing the cooperative option. If all agents have the same two options, mutual cooperation (or the mutual cooperation outcome) is all agents choosing their cooperative options. Example: if you and your friend each have the choice of doing the other a favor or not doing the other a favor, the cooperation option is doing the other a favor, and if you both choose this option, the mutual cooperation outcome (i.e., the benefits the favors provide) is the result. (p. 195)

Cooperative decisions (potentially) – any interdependent decision situation (*game*) in which *agents* do not have complete conflicts of interest, allowing for (but not guaranteeing) some *goal* achievement on all the agents' parts. Also called "nonzero sum games" because the *outcomes* of at least one pair of *options* do not sum to zero. Example: two agents who desire to go to the movies together are in a conflict about what movie to see, but are not in a complete conflict for they both desire to go to the movies together. If one or both compromise, they can both have a degree of goal achievement. (pp. 170 and 195)

Cost(s) – anything of *value* to an *agent* that the agent must use, give up, or sacrifice as a consequence of deciding on an *option*. Any cost that cannot be recovered is a *sunk cost*. Example: to drive to work, an agent will typically incur as a consequence the cost of a certain amount of wear to the vehicle driven. (p. 15)

Criterion – an *attribute*, plus its *value* or weight, in terms of which *options* and *outcomes* are compared and *ranked* or evaluated. In *single criterion decisions* there is only one category of evaluation; example: deciding which of three used cars to buy using only price as the criterion. In *multi-criteria decisions* there are at least two equal or unequal categories of evaluation; example: deciding which of three used cars to buy using price, fuel consumption, and resale value as criteria, with the first counting more than the other two. (p. 81)

Decisional complication – a *practical reasoning fallacy* due to *decisional conflict* and *indecision* in which an *agent* avoids a decision by attending to an excessive level of detail that unnecessarily complicates the *decision problem*. Example: an agent gives up making a decision about whom

not to invite to her wedding reception because she becomes overwhelmed with possible complications. (p. 34)

Decisional compromise – a *practical reasoning fallacy* due to the agent's social nature in which social pressures work to displace a *rational choice*, given the *goal*, with a more socially acceptable decision. Example: an agent whose goal is to dress warmly ends up dressing stylishly in cloths that do not keep him warm, yet has not given up the goal to dress warmly and does not have dressing stylishly as a goal. (p. 30)

Decisional conflict – a state of *indecision* in a *decision problem* created by either approach-avoidance, approach-approach, or avoidance-avoidance conflict. Decisional conflict is a category of *practical reasoning fallacies*. Example: an agent can't (yet must) decide between two options each of which has serious risk of total goal loss (avoidance-avoidance conflict). (p. 33)

Decisional habit – a *practical reasoning fallacy* due either to the agent's social nature or to *decisional conflict* in which forces such as personal habits, family traditions, cultural upbringing, or the pressure of *indecision* work to move an agent away from a *rational choice*, given the *goal*, and toward a more familiar or comfortable *option*. Example: an agent who desires to try a new flavor of ice cream can't decide between two new flavors and so chooses a third flavor that she has traditionally eaten. (pp. 31 and 33)

Decisional passing – a *practical reasoning fallacy* due to *decisional conflict* and *indecision* in which an *agent* facing a decision by-passes practical reasoning by having someone else make the decision for the agent. Example: an agent can't decide between two dinner specials on the restaurant menu, and requests the waitress to make the decision for him. (p. 33)

Decision problem – any situation having the means-end structure of an *agent*, a *goal*, and a set or menu of at least two *options* from which the agent must decide or choose (discover) which option is rationally the best means to achieve that goal. Decision problems are categorized by their structure or form into models that serve to organize specific content into concrete decision problems. Example: a high school senior who desires to go to a good college has been accepted to three colleges and must decide which one to attend. (pp. 4 and 278)

Default certainty (vs. direct evidence certainty) – in decisions under conditions of *certainty*, an *agent* might be practically certain about the *state* of the world by default: the agent has no reason to doubt that things are, and will remain, the same way they have (always or typically) been; the agent has (or knows that there is) no evidence to the contrary. In contrast, direct evidence certainty means that the agents has reason (has evidence) to be certain that the required state exists. Example (default certainty): if an option requires an agent to drive across a bridge for the goal (say, visiting a friend) to be achieved, the agent is certain that the bridge is usable by default if the agent knows that there has been no information that something happened to close that bridge. (p. 102)

Defection – in a potentially *cooperative game* in which an *agen*t has two *options*, cooperate with the other agent or don't cooperate with the other agent, defection is choosing the non-cooperative option. If all agents have the same two options, mutual defection (or the mutual defection outcome) happens if all agents choose their non-cooperative options. Example: if you and your friend each have the choice of doing the other a favor or not doing the other a favor, the defection option is not doing the other a favor, and if you both choose this option, the mutual defection outcome (i.e., loss of the benefits the favors would provide) is the result. (p. 195)

Degree of confidence – in decisions under *risk*, the strength, less than *certainty* and greater than *ignorance*, with which an *agent believes* the *state* exists (or that the statement that the state exists is true). In decisions with single-stage states, an agent's degree of confidence is reasonable if it matches the probability of the state(s) existing (or matches the strength of the evidence supporting the belief). In decisions with multi-stage states, an agent's degree of

confidence is reasonable if it matches the probability of the *conjunction* of the required states. Example: if your car has stalled 3 times out of the last 10 red lights you stopped for (thought not the last 3 consecutively), it is reasonable to be only 7/10 (.67) confident that it won't stall at the next red light you come to. (pp. 19 and 106)

Discounting the future -- a kind of *value* distortion in which an *agent* decreases the value assigned to a future *outcome* that would be given greater value if the outcome were present. Example: a college student who values a college degree, values it much less as a freshman when the degree is several years away, but then assigns it much greater value as a senior about to graduate, yet it is the same college degree. (p. 42)

Disjunction rule – a rule for combining probabilities stating that the disjunction of probabilities is calculated by addition. In decisions under *risk*, for each option the probabilities of the possible *states* should be mutually exclusive (one or the other state will exist, but not both) and jointly exhaustive (all possible states have been listed). For two possible states (a and b) of an option: P(a or b) = P(a) + P(b). Example: if one of your options involves a flip of a fair coin with heads (state a) achieving your goal and tails (state b) losing your goal, then for that option you should be certain of one or the other: P(a or b) = P(.5) + P(.5) = 1.0. (p. 114)

Disutility – negative *utility*; a predicted measure of goal loss assigned to an outcome. (pp. 14 and 92)

Dominance – one *option* (A) dominates another option (B) if in comparing the *outcomes* of A with those of B at least one outcome of A achieves more of the *goal* than any outcome of B, and no outcome of A achieves less of the goal than any outcome of B. Dominance is an important principle of *rational choice*, justifying a dominant option as a rational choice and the elimination of dominated options as *irrational choices*. Example: if your goal is to eat good pizza and one pizza place makes better pizza than another place down the street, the first dominates the second place. If your goal is to eat good pizza and drink good beer, and two pizza places are equal in quality of pizza but one has better beer, it dominates the other. (p. 124)

Equilibrium – in a *game*, two *options* (or their *outcomes*) are in equilibrium if no *agent* could improve the outcome by singly switching to another option. Equilibrium (in potentially *cooperative* games: *Nash equilibrium*) options are each the best available decision an agent can make in response to the other agent's decision. Equilibrium is an important principle of *rational choice*, justifying an option pair as decisions agents should make. Example: if you can drive on the right or on the left and everyone else is driving on the right, your option to drive on the right is in equilibrium with everyone else's decision, any one who singly switches to the left will have made an irrational choice. (p. 177)

Equilibrium selection problem – a problem in potentially *cooperative games* containing more than one *Nash equilibrium* for which there are no methods of *practical reasoning* or principles of *rational choice* for deciding on an *option*. This problem appears to represent a limit or breakdown of practical reasoning for achieving cooperation. Example: the games *clash of wills* and *chicken*, at least in their *symmetrical* version, are test cases or challenges for practical reasoning in so far as both suffer from an equilibrium selection problem. (p. 231)

Equi-probability error – the mistaken *belief* in combining initial *pure* or *factual probabilities* that there are only two possibilities for any event: it can happen or not happen, and so the certainty of 1.0 (that it must happen or not happen) must be divided equally between its happening (.5 chance) or its not happening (.5 chance); and so, for any possible event, there must be a 50/50 chance of its happening. Example: you are eating a hamburger and the thought occurs to you that you might choke on the food in your mouth, and think that there is a 50/50 chance of this happening because there are only two possibilities and each are equally possible: you choke on what you are eating or you don't choke on what you are eating. (p. 118)

Evaluation (vs. value judgment) – any method of rating or *ranking* items in the order in which they meet or satisfy *criteria* or standards. To be distinguished from value judgment, which is the expression in a judgment of the *subjective value* someone assigns to something. Example (of evaluation): rate the last three cars you have driven according to the criterion: fuel efficiency. Example (of value judgment): say which you desire (value) more, a vacation at the beach or in the mountains. (p. 37)

Event of interest – the event or property you are interested in discovering the probability of its happening. In decisions under *risk*, an *agent* is interested in the probability that the *state* of the world required for an *option* to have an *outcome* will happen or will exist. Example: if you are deciding whether or not to engage in risky sexual behavior, the event you should be interested in is the probability of getting an STD. (p. 109)

Evolutionary game theory – a behaviorist/naturalistic alternative to *rational choice* theory, based on the idea that biological evolution can be understood as a mechanism of selection or choice (a vast system of iterated *games*) for generating, under certain conditions, genetically based *cooperative* behavior as *outcomes*. The promise, being actively explored, is that evolutionary game theory might explain the origin of (human) *cooperation* for cases where rational choice theory fails to justify cooperation as the rational choice. (pp. 233 to 240)

Expected monetary value – for decisions under *risk* in which *outcomes* are assigned monetary *values* instead of *utility* values, the expected monetary value (EMV) of an *option* is the remainder of (1) the sum of the amounts of money assigned to each of its possible outcomes, after these amounts have been discounted by the agent's *degree of confidence* that the *state*(s) exist required for each outcome to happen, minus (2) the monetary *cost* of choosing that option. (p. 126)

Expected utility – in a decision under *risk* in which *outcomes* are assigned *utility values*, the expected utility of an *option* is the sum of the *utilities* (and *disutilities*) of the option's possible outcomes, after each utility or disutility has been discounted by the agent's *degree of confidence* that the *state*(s) required for the outcome to happen exists. (p. 131)

Factual probability – the probability that an *event or property of interest* happens based on the record of the frequency with which it has happened within a *sample space*. Example: if you are interested in the chance of rain tomorrow, you look at the record of all past days that were similar in weather conditions to the weather conditions today (your sample space), and count how many were followed by a day in which it rained (the frequency). The factual probability is the ratio of the frequency count to the number of items in your sample space, expressed as a decimal. (p. 109)

Fiduciary – any *agent* recognized (authorized) to be making a decision in an official or professional capacity on behalf of another individual. Example: a lawyer or a financial institution that makes decisions to promote the goals of a client. (p. 9)

Final outcome utility – for decisions under *certainty*, the final outcome utility is a measure of the amount of *goal* achievement an *option* (or its *outcome*) is predicted to give an agent that chooses that option. Final outcome utility is arrived at by summing an option's *weighted outcome utilities*. The option having the largest (maximum) final outcome utility is the *rational choice* for a decision under certainty. (p. 83)

Free ride outcome – in a potentially *cooperative game* in which an *agen*t has two *options*, cooperate with the other agent(s) or don't cooperate with the other agent(s), the free ride *outcome* results for the agent who chooses *defection* when the other agent(s) chooses *cooperation*. Example: if you and your friend each have the choice of doing the other a favor or not doing the other a favor, and you choose not to do the favor but your friend chooses to do you the favor, you receive the free rider's outcome (you receive a favor but give nothing in return) and your friend receives the *sucker's* outcome (provides a favor but gets nothing in return). (p. 196)

Gambler's fallacy – the mistaken *belief* in combining initial *pure* or *factual probabilities* that the probability of an event in a series of independent events can be influenced by events earlier in the series. Example: an agent who drives to work each day (assuming these are independent events) starts to believe that the probability of a car accident while driving to work is increasing because after years of accident-free driving she's increasingly "due" for one. (p. 117)

Game – a decision problem involving two or more *agents* in which the *rational choice* for each agent depends on which *option* the other agent(s) chooses. Games are either *competitive* (agents have complete conflicts of interests) or potentially *cooperative* (allowing some or all agents to have a degree of goal achievement), 2-person (involving two agents) or n-person (involving more than two agents). The methods of *practical reasoning* in game decisions are *strategic* (taking another agent's practical reasoning into account) and based on the assumption that agents have *common knowledge*. Example: each of two gas stations on opposite sides of the street must decide on a price to charge for a gallon of gas as a way to attract customers. (pp. 20 and 168)

Goal – anything an *agent* desires to achieve for which the agent must decide how to achieve it. Goals are wholes containing parts (*objectives*). A *simple goal* contains one part and is achieved once its part is achieved. A *complex goal* contains more than one part and is achieved by achieving all its parts. For purposes of practical reasoning, goals are very broadly defined; they might improve an agent's physical, social, or mental well-being, or not; they might be narrowly self-interested, or they might involve self-sacrifice on an agent's part to improve the well-being of others (*stakeholders*). Goals are subject to evaluation as good or bad by standards outside *rational choice* theory (e.g., by moral, religious, legal, or social norms/values) and by methods that are not (typically) part of *practical reasoning*. Example: an agent who desires to help a friend cheat on a test now has this as her goal, and must decide how best to help her friend cheat. (pp. 4, 11, and 52)

Good decision – a *rational choice*, a decision arrived at by applying the principles of *rational choice* and methods of *practical reasoning* appropriate to the *decision problem*. A good decision is not the same as a decision that has a good outcome or consequence. Example: a doctor who, by all the available evidence and proper medical practice, decides to prescribe an antibiotic to a patient makes a good decision, even if the antibiotic ends up harming the patient. (p.15)

Harmony – a potentially *cooperative game* in which each agent prefers *mutual cooperation* to a *free ride*, a free ride to being a *sucker*, and being a sucker to *mutual defection*, and so the cooperative *option* is the *rational choice* for each agent. (p.197)

Hope limit – in an uncertain *decision problem*, the best *outcome* an *agent* can hope for, given the *goal*. Example: an agent deciding between two different surgery procedures can hope for no better outcome than a complete return to good health, given that this is a possibility for at least one procedure, and that good health is the goal. (p. 104)

Ideal rational agents – perfect decision makers, not limited by *practical reasoning fallacies*, insufficient time, gaps in information, distractions, or any other real-world problems that could make practical reasoning go wrong or result in an *irrational choice*. The ideal rational agent is a device or conception used by abstract parts of *rational choice* theory to discover, investigate, and test certain formal properties of decisions (such as consistency) that contribute to establishing norms of rational choice. Ideal rational agents are contrasted with real agents (ordinary or expert decision makers, often college students who volunteer for experiments) observed and tested by concrete (empirical, behavioral) parts of rational choice theory. (p. 11)

Ignorance – a *decision problem* in which, for each *option*, the *agent* is (i) uncertain that the *state* required for the option to have its *outcome* exists (or will exist at the time an option is acted on, if chosen), and (ii) knows that he cannot assign a probability (form a *reasonable degree of*

confidence) that it exists (or will exist). Ignorance is one of three decision situations, the other two being decisions under (conditions of) *certainty*, and under *risk*. Example: while driving in an unfamiliar country, an agent becomes completely lost and must decide to go right or left at a fork in the road, but has no basis (and knows she has no basis) on which to guess which fork is the more likely way back to the hotel. (pp. 18, 151 and 187)

Ignorance condition – in the iterated (weakened) *prisoner's dilemma*, one of three *common knowledge* conditions that help to justify and sustain mutual *cooperation* as the *rational choice* and allow *agents* to avoid the *backward induction problem*: all agents are ignorant of the number of times they will be interacting. This condition is believed to work only if two other conditions are met: *retaliation* (a willingness to punish being *suckered* with a return of a *free ride*) and *reputation* (establishing a history of choosing to cooperate rather than defect). (p. 228)

Ignoring the base rate frequency – the mistaken *belief* in assigning initial *pure* or *factual probabilities* to a *state* that a rare event (ignoring that it might take place within a large or frequently occurring base rate) must be less probable than a frequent event (ignoring that it might take place within a small or rarely occurring base rate). Example: which happens more often in the United States: pizza ordered with fruit topping or a car that slides while driving on snowy roads? Pizza with fruit topping is rare among all the pizzas ordered, but there are millions of pizzas ordered, say, every weekend in the US (high base rate). Sliding while driving on snowy roads commonly happens, but most of the US has no snow and the parts that do don't have snowy roads all the time (low base rate). Answer: in the US, pizza ordered with fruit topping (say, pineapple) happens more often (no doubt far more frequently!) than the event of sliding while driving on snowy roads. (p. 112)

Indecision – inability to make a decision, typically in a situation of conflict, when a decision should be made. Indecision is to be contrasted with the *option* "not to decide for now." The decision to do nothing (the decision to take a wait-and-see position in a decision problem) is not the same thing as being indecisive. Example: if you find two equally priced equally wonderful birthday gifts for your mother, but can only afford to buy one, you are indecisive if you can't make up your mind which one to buy, but you are not indecisive if you opt not to decide between them today and instead wait until one goes on sale. (pp. 14 and 33)

Indifference – one of two relations (the other being *preference*) between pairs of *options* in a *decision problem* that rationally orders all options according to how much *goal* achievement (from most to least) each has been *evaluated* to yield. An *agent* should be indifferent between two options if (and only if) they have been evaluated to yield equal goal achievement; choosing one option over the other neither increases nor decreases the rationality of the choice, an agent is equally justified choosing either option. Indifference is to be contrasted with disinterest (*opting-out*), the latter referring to an agent no longer interested making a decision or achieving the goal. (p. 94)

Individual agent – any *agent* making an *individual decision* (as opposed to a decision by a *social agent* that contributes to a *social choice*). An individual decision is intended to promote the welfare of an individual or achieve an individual goal (as opposed to a social goal or a common good). The individual whose welfare is promoted by an individual decision might be the agent making the decision (which might be done out of self-interest), or some other individual(s) (for example, a client, a loved one, a stranger in need, or a group). Example: a business deciding which of 5 businesses to sub-contract a project is making an individual decision. (p. 9)

Individual decision theory – a part of *rational choice theory* that investigates the methods and principles *individual agents* (should) use to make individual decisions. Applied or behavioral or empirical individual decision theory investigates the concrete individual decisions of real (expert or ordinary, typically college student volunteers) individual agents, while abstract individual decision theory investigates the individual decisions of *ideal rational agents*. Normative individual decision theory provides principles and methods for justifying how decisions ought to be made,

while descriptive individual decision theory discovers the principles and methods of how (ideal or real) decisions are made. (p. 9)

Irrational choice – any *option* in an *agent's* menu of options that, after evaluation by appropriate methods of *practical reasoning* and applicable principles of *rational choice*, is found to violate those methods or principles and (typically) achieve less of the *goal* than another option in the menu. Example: if your goal is to eat great pizza (but not to save time or gas, etc.) and a pizza place 5 miles away makes better pizza than a pizza place 1 mile away, it is an irrational choice to decide on the closer place. (p. 28)

Irrelevant consequences – the sub-set of the total consequences of an *option* that do not count as *outcome* (positive or negative) and are not *costs.* Example: if you decide to go to the movies instead of staying home one day during summer, an irrelevant consequence is that your body temperature will raise the ambient temperature in the theater a small amount, making the theater's air conditioning system work a little harder. (p. 15)

Loss aversion – a *practical reasoning fallacy* due to *value* distortion in which an *agent* assigns a greater value to something lost than is assigned to gaining (or failing to gain) the same thing. Example: in negotiations over wages, an agent will tend to place more value on a $1 decrease in hourly pay (a loss) than the value the agent places on $1 gained in hourly pay (or on a $1 increase in hourly pay that was failed to be gained in the negotiations), and yet in each case it is the same $1 per hour that the agent is valuing. (p. 40)

Maximax – short for the *rational choice* rule: choose the maximum of the maxima (the best of the best, or the greatest of the goods). In decisions under *ignorance*, in which *outcomes* represent various degrees of large *goal* achievement (so-called "can't lose" decisions), maximax is a principle of rational choice stating that *agents* should choose, from among the set of each option's best outcome, the outcome having greatest goal achievement. Example: if your goal is to eat good pizza and there are three places that you think make great pizza (but can't assign probabilities that they do), choose the one that you think edges out the other two in quality of pizza. (p. 157)

Maximin – short for the *rational choice* rule: choose the maximum of the minima (the best of the worse, or the lesser of the evils). In *decision problems* under *ignorance*, in which *outcomes* represents various degrees of harm or *goal* loss (so-called "painful choices"), and in *games*, maximin is a principle of rational choice stating that *agents* should choose, from among the set of each *option's* worse outcome, the option whose outcome has least goal loss. Example: if your goal is to be an honest person, and you are forced to choose between lying to your boss or lying to your co-worker (and can't assign probabilities to possible outcomes), the maximin principle says that it is more rational to choose the lie that you think makes you least dishonest. (pp. 154 and 182)

Multi-criteria decisions – *decisions problems* in which *goals* are complex and require a method of *evaluating* the goal achievement of *options* with respect to multiple *objectives* in order to discover which option, on balance, has maximum outcome *utility*. Example: if your goal is to eat a dessert that is rich in chocolate, nutritious, delicious, non-fattening, and inexpensive, you have a multi-criteria decision problem requiring you to evaluate possible desserts by a balance of these 4 criteria, for you will not find one dessert that satisfies all 4 to the maximum degree. (p. 87)

Negotiation – methods of *practical reasoning* for transforming a potentially cooperative *game* containing an *equilibrium* selection or a (Pareto) *sub-optimal* problem into a *bargaining* problem. Negotiations must settle difficult problems such as agents agreeing on their independent worth (relative to the goal) and agreeing on enforcement of agreements. "Good faith" negotiation is based on enough progress to show that agents are not just engaged in stalling tactics or manipulating the process; "serious" negotiation is based on progress in the more difficult areas requiring agreement. (pp. 246 and 253)

Nash equilibrium – in potentially *cooperative games*, any pair of *options* (or *outcomes*) that are in *equilibrium*. Named for the mathematician and Noble Prize winner John Nash who proved that every 2-person game contains at least one equilibrium outcome in pure or mixed strategy. (p. 196)

Normalizing – a form of standardizing. In a *decision problem* in which interval *ranking* is used, normalizing is a method of bringing rankings into alignment so that each ranking sums to the same unit: 1.0. Normalizing is achieves by summing the values in an interval ranking and then dividing each value that was summed by their sum. Example: if you ranked, on a scale of 1 to 10, four pizza places (a, b, c, d) according how good a pizza each makes, and got: a = 2, b = 10, c = 3, and d = 7, you normalize this ranking by summing (= 22) and dividing each value by 22. The new values sum to 1.0. (p. 62)

Objectives – parts of a *goal*. By achieving objectives, an *agent* achieves a goal. A *complex goal* has more than one objective, requiring an agent to rank or prioritize objectives in their order of importance. Objectives are to be distinguished from a "to do" list of activities related to the goal. Example: an agent whose goal is to earn a college degree has a complex goal that is achieved by achieving a set of objectives that will include: gaining admission to a college, passing required courses, selecting a major, etc. A "to do" list such as: requesting a college catalog, calling advisors, filling out financial forms, buying books, etc. is not a list of parts of the goal, and so are not objectives. (pp. 12 and 57)

Optimal mixed strategy – in a *game* in which there is no *pure strategy* (no single option is a rational choice solution), a proportion of switching among *options*, done under conditions of strategic (intentional) *ignorance*, which yields an *agent* the most amount of *goal* achievement possible in response to the decision(s) of the other agent(s). An optimal mixed strategy game is fair if agents playing their optimal mixed strategy have equal goal achievement, and biased if their goal achievement is unequal. Example: the child's game paper-scissors-stone is a fair optimal mixed strategy game, but a roulette wheel at a casino that slightly favors the house over the customer is a mixed strategy game biased against the customer. (p. 188)

Options – a set (menu) of at least two alternative *action-state-outcome* trios, each possible for an *agent* to do, from which an *agent* must choose one as the means to achieve a *goal*. Also called: choices. Example: if your goal is to go to the movies, your options (choices) are the movies playing in theaters you can get to. (pp. 12 and 52)

Option not to act – the *option* in an *agent's* menu of options to hold off on making a decision; the option to take a "wait-and-see" attitude toward the decision problem, sometimes described as "keeping our options open," other times describes as "*procrastination*." The option not to decide is to be contrasted with *opting-out*. Example: your summer plans requires a plane trip, and your options include 3 airline companies from which you could buy your ticket, but you decide to hold off making a purchase in order to wait and see if any airline will lower its prices. (p. 14)

Opting-out – the *option* in an *agent's* menu of options to give up the *goal*. Opting-out is to be contrasted with the *option not to decide* or not to act for now. Example: if you give up the goal of a college degree you are opting-out of the decision problem of how best to achieve this goal, not just holding off making a decision about the best way to achieve it. (p. 14)

Ordinal ranking value – in decisions under conditions of *ignorance*, for each *option* the *agent* has enough information about *states* (most to least probable) and *outcomes* (most to least goal achievement) to rank them ordinally. The ordinal ranking value (ORV) of an option is the sum of the products of the ordinal number assigned to each state times the ordinal number assigned to its outcome. The rational choice rule for decision problems in which ORV's can be discovered for its options is: choose the option with the largest ORV. (p. 159)

Outcome – the sub-set of the predicted consequences of doing an action that either gains (positive outcome) or loses (negative outcome) an *agent* part of or the entire *goal*. The amount of goal achievement an outcome is predicted to yield is its *utility* and the predicted amount of goal loss is its *disutility*. Example: if your goal is to visit a friend and you decide to drive, then of all the consequences that might result from driving (e.g. pollution, fuel consumption, etc.) the (positive) outcome is the consequence of arriving at your friend's location. (pp. 14, 83 and 103)

Pareto sub-optimal (optimal) – in potentially *cooperative games*, an *outcome* pair is Pareto sub-optimal if there is another available outcome pair that gives at least one *agent* more goal achievement and no agent less goal achievement. If there is no such available pair of outcomes, an outcome is Pareto optimal. Named after the Italian economist and social theorist Vilfredo Pareto (1848–1923). Example: in the prisoner's dilemma, the rational choice, mutual defection, is troubling because its outcome is sub-optimal compared to the available mutual cooperation outcome. (p. 218)

Payoff – an *option*'s *outcome* or amount of goal achievement (or loss) in a *game*. (p. 168)

Player – an *agent* whose *decision problem* is a *game*. (p. 168)

Practical reasoning – reasoning about what to do; rational decision making. A form of means-end (instrumental) reasoning that, if done well, yields a *good decision* (even if the *agent* fails to achieve the goal). Reasoning that applies methods of *framing* (analysis) and principles of evaluation within the structure of a *decision problem* in order to discover and justify a *rational choice* solution. Practical reasoning is to be contrasted with critical reasoning, which operates within an inference structure (argument) to rationally justify a belief (conclusion) on the basis of evidence (premises). Example: the deliberations you might go through to decide on a good birthday gift for your best friend. (pp. 3 and 277)

Practical reasoning fallacy – one of a number of commonly found patterns of poor decision making that results in an *irrational choice*. Example: procrastination (putting off until later an action that should be started now in order for the goal to be achieved) is an irrational choice that typically results from poor practical reasoning. (p. 28)

Preference – one of two relations (the other being *indifference*) between pairs of *options* in a *decision problem* that rationally orders all options according to how much *goal* achievement (from most to least) each option has been evaluated to yield. If one option should be preferred to another, it is more rational to choose it over the other. A preference ordering of the options is a *rational choice* solution to a decision problem. A preference order over options is normative, representing how an agent should choose or how an *ideal rational agent* would choose. (p. 93)

Prisoner's dilemma – a potentially *cooperative game* in which each *agent* prefers a *free ride* to *mutual cooperation*, mutual cooperation to *mutual defection*, and mutual defection to being a *sucker*. The prisoner's dilemma suffers from a *sub-optimal outcome* problem, the *rational choice* results in the mutual defection *outcome* yet each agent could have better *goal* achievement with mutual cooperation. The prisoner's dilemma might be a one-time decision problem or a series of repeated (iterated weakened) decision problems, in the latter case subject to the *backward induction problem*. (p. 223)

Procrastination – the *practical reasoning fallacy*, due either to *decisional conflict* or to *value* distortion, of deciding to hold off making a decision or doing an action that the agent knows must be made, or that will require a certain amount of time to achieve an agent's goal, until there is not enough time left for the decision or action to achieve that goal. Example: an agent knows she needs a full week of piano practice to achieve her goal of a successful piano recital, and can put in the required practice, but decides not to practice until two days before the recital. (pp. 33 and 44)

Pure probability – the probability of a property or *event of interest* happening bases on the number of possible ways it can happen within a *sample space* of abstract possibilities. Example: there are only 52 possible cards you can pick from a normal deck of cards, and there are only 4 possible aces, so the pure probability of picking an ace from well shuffled normal deck of cards is 4 out of 52 or 1/13. (p. 110)

Random decision – a *fallacy of practical reasoning* due to *decisional conflict* and *indecision* in which an *agent* lets chance, a random event, be the deciding factor about which *option* to choose. The fallacy of random decision should be distinguished from *strategic ignorance*. Example: it is getting late and an agent who can't decide what to wear to an important interview gets over his indecision by closing his eyes and randomly selecting the first item he grasps from his clothing rack. (p. 33)

Ranking – methods of *practical reasoning* that orders items according to a category, or a set of norms or standards (for example: an ordering of items according to size, or price, or importance, or preference). A qualitative ranking requires brief descriptions (words or phrases) to represent the rank of each item (p. 59). An ordinal ranking uses the ordinal numbers to represent rank (p. 60). An interval ranking uses cardinal numbers to represent rank as well as the size of relative distances or intervals between any two items (p.61). Example of an interval ranking: on a scale of 1 to 15, rank the last five movies you have seen in the order of how much you enjoyed the music in each. (p. 59)

Rash decision – a *fallacy of practical reasoning* due to *decisional conflict* and *indecision* in which an agent makes a hasty or quick decision just to get it over with, without the effort of practical reasoning that would justify the decision as a rational choice. Example: an agent makes a hasty decision about which medical procedure to undergo after a period of indecision worrying about possible unpleasant side-effects of each of her options. (p. 32)

Rational choice – the solution to a *decision problem* arrived at by methods of *practical reasoning* and justified by principles and norms established by the theory of rational choice. A rational choice is typically a *preference/indifference* order imposed on an *agent's options* after *evaluating* possible *outcomes* with respect to the *goal*. Example: in a decision under certainty, an agent should prefer the option having maximum final outcome utility and least prefer the option having minimum final outcome utility. (pp. 21 and 76)

Raw mean (error of) – the mistaken *belief* in assigning initial *pure* or *factual probabilities* to a *state* that the closer an event or property is to the average (the raw mean) the more probable it must be. Within a range of values, the average value might not happen at all and might even be impossible. Example: in 10 families each having a different number of children between 1 and 10, the average family size is 5.5 children; but it is impossible, and thus not probable at all, for any of these 10 families (or any family) to have 5.5 children. And it is no more probable in these 10 families to have 5 children than it is to have 1 child or 10 children. (p.113)

Relative value distortion – a *fallacy of practical reasoning* in which an *agent* enlarges or belittles the value assigned to a unit of something the agent values, relative to other units, in inconsistent ways. Example: an agent whose goal is making money refuses to accept an offer of $15 for an item he is asking $20 for at a garage sale, but then decides to accept an offer of $45 for an item he is asking $50 for (when all relevant variables make these equal decision problems). In each case there is $5 at stake, so these should be the same (consistent) decisions: refuse both or accept both. (p. 40)

Reputation condition – in the iterated (weakened) *prisoner's dilemma*, it is believed that mutual *cooperation* is the *rational choice* if three *common knowledge* conditions are met: (i) each *agent* establishes a reputation with the other of choosing cooperation over *defection*, (ii) *retaliation* (being suckered will be punished with the free ride), and (iii) *ignorance* (each agent does not

know the number of repeated interactions that will take place between them (thereby avoiding the *backward induction problem*)). (p. 226)

Retaliation condition – in the iterated (weakened) *prisoner's dilemma*, one of three *common knowledge* conditions that help to justify and sustain mutual *cooperation* as the *rational choice*: a willingness to retaliate (get even) for being *suckered* by choosing to defect (trying for a free ride) in a later iteration. This condition is believed to work only if two other conditions are met: a *reputation* for cooperation and *ignorance* of the number of iterations. (p. 227)

Risk – a *decision problem* in which, for each *option*, the *agent* is (i) uncertain that the *state* required for the option to have its *outcome* exists (or will exist at the time an option is acted on, if chosen), but (ii) can form a reasonable *degree of confidence* that it exists (or will exist as needed). Risk is one of three decision situations, the other two being decisions under (conditions of) *certainty*, and under *ignorance*. Example: while driving in an unfamiliar country, an agent becomes completely lost and must decide to go right or left at a fork in the road, but is more confident that the left fork is the way back to the hotel (say, .8 confident it's the left fork and only .2 confident the right fork is the way back) based perhaps on a vague memory of looking at a map of the area weeks ago. (pp. 18, 103 and 147)

Risk aversion – one of three possible attitudes an *agent* can have toward *risk* in decisions under risk. An agent is risk averse (and certainty seeking) if an agent prefers a certain option to a risky option when the rational choice should be *indifference*; the stronger the *preference*, the more risk averse the agent is. An agent is risk seeking (and certainty averse) if the agent prefers a risky option to a certain option when the rational choice should be indifference; the stronger the preference, the more risk seeking the agent is. An agent is risk neutral (and certainty neutral) if the agent is indifferent between a risky option and a certain option when the rational choice should be indifference. Example (risk aversion): if you prefer a sure $10 to an even chance, say by a flip of a fair coin, to gain $20, then you are risk averse, because by EMV the rational choice should be indifference between these two options. (p. 148)

Saddle point – in a 2-person game solved by *maximin* reasoning, the one cell containing each agent's maximin outcomes. (p. 153)

Sample space – in discovering the probability that a property or *event of interest* will happen, the sample space is the set of properties or events within which the event of interest happens. Example: if you are interested in the probability that an item bought will be returned by the customer, your sample space is all cases of that item being bought (or a good sample in case the number is too large), within which the event of interest (being returned by a customer) happens. (p. 109)

Satisficing – methods of *practical reasoning* within *bounded rationality* in which an *agent* forms realistic criteria of what counts as a good enough menu of *options* and *outcomes*, and makes a decision according to these criteria, rather than applying the standards of *rational choice* for *ideal rational agents*. Example: a person in a rush picks the first daily special the waiter mentions, given it is not a dish the agent dislikes, rather than letting the waiter mention all the daily specials and deciding on the best. (p. 162)

Scale – a range of values that can be used to represent an ordering, listing, or *ranking* of a set of objects. A scale might be a qualitative representation (e.g. list those 5 red sweaters in order of darkest to lightest values of the color red) or a quantitative representative (e.g., on a thermometer, a range of values that are possible human body temperatures from coolest to hottest). A scale might be a verbal representation (using words) or numerical (using numbers) such as an ordinal or an interval representation, among other possibilities. In *practical reasoning*, methods of representing values and methods of evaluations use scales, especially ordinal ranking and interval ranking scales. (pp. 60-62)

Scenario thinking – the mistaken *belief* in assigning initial *pure* or *factual probabilities* to a *state* that an accurate estimate of how probable something is can be discovered by imaginatively creating a scenario in which it happens. If an agent does not know the correct probability of an event happening, the agent will not be able to create a mental scenario that contains the correct probability of that event happening (or recognize that a probability is the correct one). Example: an agent who tries to discover the probability of her suffering food poisoning from eating at a fast food restaurant (something she does not know) by imaginatively picturing herself eating at fast food restaurants and getting food poisoning at some of these. This count will almost surely not yield an accurate probability. (p. 112)

Security level – in an uncertain *decision problem*, the *outcome* an *agent* can be assured of than which there is no worse the agent need fear. In *bargaining problems*, a level of independent *goal* achievement or loss that establishes an agent's initial bargaining position or bargaining power. Example: if a thief presents an agent with a choice "your money or your life", the agent is assure that there is no worse outcome in this decision problem than losing his life, assuming staying alive is one of the agent's goal. (pp. 104 and 247)

Simple goal – a *goal* having just one part, making it achievable by achieving its one *objective*. A simple goal is not necessarily an easily achieved goal. Example: a person who is standing up has a goal to remain standing until counted, and achieves this goal by achieving one objective: to continue standing until counted (no matter how physically hard or uncomfortable it might be for this person to remain in a standing position). (p. 12)

Single-criterion decisions – *decision problems* with *simple goals*, the goal itself serving as the single criterion for evaluating which option has an outcome that most achieves that goal. Example: If your goal is to eat a meal that contains minimal meat, knowing the meat content of the possible meals you could eat allows you to use this single criterion to select the one that is most lacks meat. (p. 70)

Social agent – any *agent* making a decision as a member of a group that is intended to contribute to a group decision or social choice. A social choice is a decision relative to a *common interest* intended to promote the welfare of a group or achieve a social goal (a common good) by accounting for or aggregating the decisions of its members. Example: citizens of a town who decide by voting who the next town mayor will be are social agents. (p. 9)

Social choice theory – a part of *rational choice theory* that investigates principles and methods of arriving at social choices on the basis of decisions of *social agents*. There are major challenges in the alternative methods of aggregating the decisions of social agents into a social choice, such as the voting paradox, in addition to the problem that different methods of aggregation can yield different social choices. (p. 10)

Stable – a *game* is stable if the methods of *practical reasoning* bring each *agent* independently to choose a single *option* as the solution, from which no agent has reason to switch to another option. (p. 177)

Stag hunt – a *potentially cooperative game* in which each *agent* prefers *mutual cooperation* to a *free ride*, a free ride to *mutual defection*, and mutual defection to being a *sucker*. The stag hunt suffers from a *sub-optimal outcome problem* in so far as mutual defection is more attractive (a sub-optimal *maximin equilibrium outcome*) than mutual cooperation (an *optimal* equilibrium outcome involving *risk* and requiring trust). (p. 217)

Stakeholder – any individual(s), e.g. person(s), organization(s), group(s), a decision is intended to benefit. Stakeholders are to be distinguished from any individuals that happen to be benefited or otherwise affected by a decision as an *unintended consequence*; the former as such are part of the *goal* while the latter are not. Example: if I buy you a gift for your birthday, my decision what

to buy has you as a stakeholder, but the shop owner who profits from my purchase is not a stakeholder. (p. 19)

State – the state-of-the world or relevant conditions that must exist if an action that an *agent* might choose to do is to result in the desired *outcome*. An agent's *belief* concerning the existence of the state each *option* requires might be (i) *certain*, (ii) a degree of uncertainty (*risk*), or (iii) *ignorance*. The state might be single-stage (a single required situation or condition), or it might be multi-stage (a series of required situations or conditions). Example: deciding to drive to work, instead of taking public transportation, will result in your goal of arriving at work provided a working car, traffic conditions, and other relevant conditions exist. (pp. 15 and 107)

Strategy – an *option-outcome* pair in a *game*. (p. 168)

Strategic ignorance – in *mixed strategy games*, an intentional form of ignorance on an *agent's* part about which *option* will be chosen so as not to distribute choices among options in a predicable pattern that could be discovered by other agents. Example: an agent that must choose among 3 options in strategic ignorance might roll a fair die and choose option 1 if side 2 or 5 comes up, option 2 if sides 3 or 6 lands up, and option 3 if sides 1 or 4 lands up; no one would or could know beforehand, including the agent, which option the agent will decide on. (p. 187)

Strategic reasoning -- *practical reasoning* in a *game* requiring each *agent* to take into account the practical reasoning of other *players*. An important part of strategic reasoning is the *common knowledge* assumption. Example: if you will decide to go to a party only if someone else decides not to go to that party, your decision to go or not must try to take into account what that other person will decide about going to the party. (pp. 20 and 169)

Subjective value – the amount or degree to which something is desired, as represented on a *scale*. To be contrasted with objective value, the amount or degree of good (goodness) something has as a property or status independent of anyone's desire for it. Example (subjective value): if a person desires a college degree twice as much while a college senior as she did while a college freshman, then the degree has twice as much subjective value for her as a senior than it did when she was a freshman. (p. 37)

Sub-optimal outcome problem – a problem in some potentially *cooperative games* in which the *rational choice* leaves *agents* with less *goal* achievement than they could gain by making *irrational choices* (or less rational choices), which are often the *mutual cooperation* options. There are no (or out-numbered) methods of *practical reasoning* or principles of *rational choice* justifying as rational the (cooperative) *option* yielding the optimal *outcome*. This problem appears to represent a limit or breakdown of practical reasoning for achieving cooperation. Example: the games of stag hunt and prisoner's dilemma are test cases or challenges for practical reasoning in so far as both suffer from the sub-optimal outcome problem. (p. 232)

Sucker's outcome – in a potentially cooperative *game* in which an *agent* has two *options*, cooperate with the other agent or don't cooperate with the other agent, the sucker's outcome results for the agent who chooses to cooperate when the other agent chooses *defection*. Example: if you and your friend each have the choice of doing the other a favor or not doing the other a favor, and you choose to do the favor but your friend chooses not to do you the favor, you receive the sucker's outcome (you give but get nothing in return). (p 196)

Sunk cost – anything an *agent values* that must be invested or used in order for an *option* to have an *outcome* but which can't be recovered. Sunk costs can trap an agent in a decision that it is no longer rational for the agent to make (practical fallacy of sunk costs). Example: a decision requiring an agent to invest a large amount of time that can't be recovered or made up has the value of this time as a sunk cost. (pp. 15 and 35)

Symmetrical game – any *game* in which it would not matter to the *agents* which *option-outcome* position each would be assigned; given the options, the level of goal achievement would be the same no matter which position an agent is assigned in the game. To be contrasted with *asymmetrical games*. Example: two agents are flipping pennies with heads winning a nickel and tails losing a nickel would be (or should be) indifferent as to which flips first and which flips second. (pp. 177 and 246)

Uniqueness error – the mistaken *belief* in assigning initial *pure* or *factual probabilities* to a *state* that the more unique a property or event is the less probable it must be. Example: in picking two cards from a normal well-shuffled deck of cards, which is more probable: picking a king and a queen or picking, say, a 5 and a 7? The first possibility might seem more unique because there are fewer kings or queens in a normal deck of cards than there are numbered cards; but the probability of picking each pair is exactly the same. (p.113)

Unintended consequences – a *practical reasoning fallacy* due to the *value* distortion of *discounting the future* in which an *agent* reduces or neglects the negative value of *costs*, *outcomes*, or consequences of an *option* that will occur in the relatively distant future. Example: an agent who has decided to start smoking as a teenager reduces, fails to take seriously, or doesn't even consider the negative value to the agent of the consequences of poor health and losing several years of her life because they won't happen until decades later.
(p. 43)

Utility – for each *outcome*, the measure of the amount or degree or strength with which it achieves (or loses) the *goal*. If an outcome is, or brings about, a degree of goal achievement it has positive utility; but if it is, or results in, any amount of goal loss it has negative utility or disutility. A central purpose of methods of analysis and *evaluation* in *practical reasoning* is discovering and assigning utility to outcomes, and several principles/norms of rational choice are expressed in terms of utility. Example: if your only goal is money and you have to choose between 2 job offers, and one job has a salary twice as large as the other, then the outcome of taking it should have twice as much utility as the outcome of taking the other job. (pp. 14 and 92)

Value (relative) – *subjective value* that varies relative to changes in something else. If the subjective value something has for an agent is independent of any changes taking place in other things, it is intrinsic subjective value. Examples (relative value): the more trouble a person's car gives her, the less she values it; the less money a person has, the more he values his job. Example (intrinsic value): the more trouble the new baby is to the new parents, their love (value) for their child remains the same. (p. 38)

Value distortion – a category of *practical reasoning fallacy*. Value distortions are inconsistencies in *valuing* items in similar or basically identical *decision problems* due to the influence of differences in the decision situations that, upon analysis and *evaluation*, are irrelevant. Value distortion violates the principle: valued items in decision problems that are equal in all relevant respects should be valued equally. Example: *discounting the future* is a value distortion due to time that can lead to the practical reasoning fallacy of *procrastination*. (p. 37)

Weighted outcome utility – for decisions under *certainty*, a *utility value* assigned to an *outcome* according to the degree the outcome satisfies a *criterion* in comparison to other outcomes, and then discounted by the weight or value of the criterion. Example: if you are comparing three used cars according to the criterion "fuel efficiency" which is to weigh heavily in your decision (say it counts 75% of your goal to buy a good used car), you would rank each car according to its average fuel use with a utility number, and then discount each utility by .75 to arrive at the weighted outcome utility for each car under that criterion. (p. 83)