

Decisions in Branching Time

Paul Bartha

Abstract This chapter extends the deontic logic of Horty (*Agency and deontic logic*, 2001) in the direction of decision theory. Horty's deontic operator, the dominance ought, incorporates many concepts central to decision theory: acts, causal independence, utilities and dominance reasoning. The decision theory associated with dominance reasoning, however, is relatively weak. This chapter suggests that deontic logic can usefully be viewed as *proto-decision theory*: it provides clear foundations and a logical framework for developing norms of decision of varying strength. Within Horty's framework, deontic operators stronger than the dominance ought are defined for decisions under ignorance, decisions under risk, and two-person zero-sum games.

1 Introduction: Decision Theory and Deontic Logic

Consider the following two decision problems.

Example 1 (*Gambler*): An agent, α , is offered a gamble. If she accepts, she pays \$5. A coin is then tossed: on *Heads* she wins \$10; on *Tails* she wins nothing. If she declines the gamble, she simply keeps her \$5.

Example 2 (*Matching Pennies*): Two agents, α and β , simultaneously choose whether to display a penny *Heads up* or *Tails up*. If the displayed sides of the two pennies match, then α wins \$1 from β . If the two sides do not match, then β wins \$1 from α .

Assuming that these agents value money positively and that there are no relevant external considerations, what should α do in these two scenarios? To find answers, we

P. Bartha (✉)
Department of Philosophy, University of British Columbia, 1866 East Mall, E370,
Vancouver, BC V6T 1Z1, Canada
e-mail: paul.bartha@ubc.ca

might look to two distinct normative frameworks: decision theory¹ and deontic logic. Decision theory, despite its many paradoxes and controversies, provides our most successful formal account of rational choice. Deontic logic, with its own paradoxes and controversies, offers an alternative way to think about what α ought to do.

What is the relationship between decision theory and deontic logic? We might think of them as directed towards answering different questions. Decision theory rests on sharp assumptions about preferences and belief, but also upon less clear assumptions about causation, choice and counterfactuals. Deontic logic has no place for probabilities or probabilistic reasoning and does not pretend to offer a comprehensive theory of rational choice. Yet both theories can be applied to examples like *Gambler* and *Matching Pennies*. This suggests that they might, in some way, be rivals.

There is a third possibility. Rather than see them as unrelated or as rivals, we might regard deontic logic as a kind of *proto-decision theory*. Conceived in this way, deontic logic would play three roles. First, it would provide a rigorous logical framework for decision theory, a framework that clarifies foundational assumptions about causation, choice, counterfactuals and other relevant concepts. Second, stronger and weaker systems of deontic logic would be definable in this common framework. Third and finally, these systems of deontic logic would also be rudimentary decision theories: they would provide norms for choices by agents that are compatible with basic principles of decision theory.

When deontic logic is viewed in this way, the approach developed by Horty (2001) is exemplary. Horty's deontic logic (in contrast to many earlier approaches) is *prescriptive*: it is about choices by agents. It proposes semantics for what a particular agent *ought to do* at a particular moment in time. Horty's framework is built on top of Belnap's modal logic of agency, *stit* theory, a clear and rigorous logical and metaphysical account of agents making choices in indeterministic branching time.² *Stit* theory already takes us part way to causal decision theory because it incorporates causal notions: agents, branching time and a formulation of causal independence. Horty takes us further by incorporating utilities and dominance reasoning into his account, although he does not introduce probabilistic concepts. Still, his deontic logic does provide a weak decision theory: under reasonable assumptions, the set of obligations for an agent on the Horty semantics is a subset of the set of obligations that the agent has according to any sound principle of decision theory.

The main thesis of this chapter is that Horty's approach can be fruitfully enriched, first by a slight generalization of his account of causal independence and second by adding a 'thin' layer of probabilistic concepts.³ The result is a framework in which deontic logic is even better suited to serve as *proto-decision theory* by playing the

¹ By 'decision theory' I mean to include the theory of decisions under ignorance, decisions under risk and normative game theory.

² This account is developed by Belnap and others in a series of articles, many of which are reprinted in (Belnap et al. 2001).

³ In a similar spirit, Kooi and Tamminga (2008) show how Horty's framework can be supplemented to engage with game theory (though without introducing probabilistic ideas).

three roles mentioned above. There may be good reasons not to introduce full-blooded probability into the *stit* universe.⁴ But it is plausible to introduce probabilities for mixed strategies by agents, and more generally for chance mechanisms (such as coin tosses and dice rolls). Horty's deontic logic can then be expanded fruitfully towards different branches of decision theory.

The chapter proceeds as follows. Section 2 reviews the basic ideas of *stit* (seeing-to-it-that). Section 3 proposes a generalization of Horty's account of causal independence. Horty's *Dominance Ought* is reviewed in Sect. 4, along with a slight modification corresponding to the generalization of Sect. 3. The remaining sections explore expansions of Horty's deontic framework to decisions under ignorance (Sect. 5), decisions under risk (Sect. 6) and elementary game theory (Sect. 7). While the focus of the chapter is on endowing deontic logic with resources from decision theory, I conclude with a brief discussion of return benefits for decision theory.

2 Seeing to it That (*stit*)

In a series of articles, Belnap and others have proposed semantics for the modal construction, “ α sees to it that A ,” or $[\alpha \textit{ stit}: A]$ for short. To keep things brief, I pass over the philosophical motivation and simply review concepts that are crucial for this chapter. The best single source of information on *stit* is the volume of articles (Belnap et al. 2001).


2.1 Semantics for *cstit* with One Agent

There are three accounts of $[\alpha \textit{ stit}: A]$ on offer: the Belnap “achievement *stit*” (*astit*), the Horty/von Kutschera “deliberative *stit*” (*dstit*) and the “Chellas *stit*” (*cstit*). Since the latter is employed by Horty in his deontic logic, this section presents, in cursory form, only the semantics for the Chellas *stit*, following notation that borrows from both (Belnap et al. 2001) and (Horty 2001).⁵ The fundamental idea of $[\alpha \textit{ cstit}: A]$ is that A is guaranteed by a present choice of agent α ($[\alpha \textit{ dstit}: A]$ is more complex because it requires this same *positive condition*, together with the *negative condition* that A is not ‘settled-true’ in the sense of (3) below).

The framework begins with an indeterministic branching time structure $\langle Tree, <, > \rangle$, where *Tree* is a non-empty set of *moments*, m , and $<$ is a tree-like partial ordering of those moments. A *history*, h , in *Tree* is a maximal chain of moments, i.e., a complete temporal evolution of the world. If m is a moment, write

⁴ Belnap et al. (2001) consider and reject the idea that “ α sees to it that A ” should be modelled as “ α ’s choice guarantees a high probability for A ”; Broersen develops this very notion in (2011) and elsewhere. In this chapter, I am concerned not with a probabilistic version of *stit* but with the importance of probabilities for norms of choice.

⁵ Both *dstit* and *cstit* have been useful in deontic logic (Belnap et al. 2001). On *dstit* versus *cstit* as an analysis of seeing-to-it-that, see (Chellas 1992) and (Horty 2001).

Fig. 1 Histories

$H_m = \{h/m \in h\}$ for the set of all histories containing (passing through) m . The situation is illustrated in Fig. 1, where the upward direction represents later moments.

In the picture, $H_m = \{h_1, \dots, h_6\}$, while $H_{m_2} = \{h_1, h_2, h_3\}$. The histories h_1 and h_2 are *undivided* at m because they share a later moment (m_2) in common; the histories h_1 and h_4 are *divided* at m .

Sentences are constructed from propositional variables A, B, \dots using the following operators:

- (i) Truth-functional operators: \sim, \vee (with abbreviations \wedge, \supset, \equiv)
- (ii) Necessity operators: *Universally*:, *Settled*:
- (iii) Tense operators: *Will*: and *Was*:
- (vi) Agentive operator: $[\alpha \text{ } cstit: \text{ } __]$ where α denotes an individual agent

The truth of sentences is evaluated relative to a moment-history pair m/h , where m is a moment belonging to history h . A *model* M pairs the tree structure with an interpretation that maps each propositional variable A to a set of m/h pairs where A is true:

- (1) $M, m/h \models A$ iff A is true at m/h .

In Fig. 1, for example, $M, m_2/h_1 \models A$ while $M, m_3/h_4 \not\models A$. The clauses for the truth-functional operators are standard. For the modal operators, *Universally*: represents truth throughout *Tree*, while *Settled*: represents truth throughout a moment. The relevant clauses are as follows:

- (2) $M, m/h \models \text{Universally: } A$ iff $M, m'/h' \models A$ for all m'/h' in *Tree*
- (3) $M, m/h \models \text{Settled: } A$ iff $M, m/h' \models A$ for all h' with $m \in h'$

For the tense operators, we have

- (4) $M, m/h \models \text{Will: } A$ iff $M, m'/h \models A$ for some m' in h with $m < m'$
- (5) $M, m/h \models \text{Was: } A$ iff $M, m'/h \models A$ for some $m' < m$.

In Fig. 1, *Will: A* is true at m/h_1 but false at m/h_4 . *Settled: A* is true at m_2/h_1 but false at m_3/h_5 .

Finally, we come to $[\alpha \text{ cstit}: A]$. This requires enriching the branching time framework of $\langle Tree, < \rangle$ with a nonempty set $AGENT$ of agents (denoted α, β , and so forth) and a function $Choice$ that represents choices by agents. The most important idea is that of a *choice set* for α at moment m , which is a partition of the histories passing through m into *choice cells* (or simply *choices*) for α . α 's power of choice consists in "constraining the course of events to lie within some definite subset of the possible histories still available". (Belnap et al. 2001, 33). That is, choice is identified with the selection of one cluster of histories. The agent picks the cluster, but cannot select a unique history within the choice cell. All of this is formalized in the following definitions.

(6) *stit frames*.

A *stit frame* $\langle Tree, <, AGENT, Choice \rangle$ is a structure with $Tree$ and $<$ as above, $AGENT$ a nonempty set of agents, and $Choice$ a function mapping agent α and moment m into a partition of H_m characterized as follows:

- $Choice_\alpha^m$ is a partition of H_m into mutually exclusive and exhaustive sets.
- Each member of $Choice_\alpha^m$ is called a *choice cell* (or *choice*) for α at m
- h and h' are *choice-equivalent* for α at m (written $h' \equiv_m^\alpha h$) if they belong to the same choice cell for α at m (no choice that α can make at m tells them apart).

$Choice$ is subject to the following condition (and one further condition, *Weak Independence of Agents*, to be described shortly):

(7) *No Choice between Undivided Histories*.



If h and h' are undivided at m , then h and h' must belong to the same choice cell for α at m .

A choice by some agent is the most obvious means by which histories are divided. Histories also divide as the result of chance processes in *Nature*. A coin toss serves as a paradigm example.

With this apparatus on board, we can define what it is for $[\alpha \text{ cstit}: A]$ to hold at m/h :

- (8) $M, m/h \models [\alpha \text{ cstit}: A]$ iff A is true at (m, h') for all h' with $h' \equiv_m^\alpha h$. (By choosing the cell containing h , α guarantees that A is true as A holds on all histories consistent with α 's choice.)

Figure 2 provides the basic picture. In the picture, m is a moment in $Tree$ that has been blown up to reveal the choice structure. The three boxes represent choice cells for α at m ; each history in H_m belongs to exactly one box. The truth-value of A is shown for each moment–history pair. Here, $[\alpha \text{ cstit}: A]$ holds just at m/h_1 and m/h_2 (Note that $[\alpha \text{ cstit}: \sim A]$ is false throughout m ; there is no *law of excluded middle* for seeing-to-it-that).

Fig. 2 [α *cstit*: A]**Fig. 3** Multiple agents

2.2 Multiple Agents: Independence and Joint Agency

The concept of *cstit* generalizes to groups of agents. Most of the ideas can be made clear by considering just two agents, α and β , making simultaneous choices. The simplest way to represent this is once again with a blown-up picture of the moment, this time two-dimensional, with the choices of α represented on the horizontal axis and the choices of β on the vertical axis, as in Fig. 3.

Here, α and β face non-trivial choices, formally specified by choice sets (partitions) $Choice_\alpha^m$ and $Choice_\beta^m$. In the picture, the possible choices are a_1 and a_2 for α , and b_1 and b_2 for β ; thus, the choice sets are $\{a_1, a_2\}$ and $\{b_1, b_2\}$.

The key assumption made by Belnap and Perloff is that every combination of choices by α and β is possible:

(9) [Weak] *Independence of Agents*.

For each moment and for each way of selecting one choice for every agent (in the set *AGENT*) from among that agent's set of possible choices at that moment, the intersection of all the choices selected must contain at least one history.

To formalize this condition, we follow Horty in defining a *selection function* s at moment m to be a mapping from *AGENT* into H_m that selects one action for each

agent: $s(\alpha) \in \text{Choice}_\alpha^m$ for each α . Let Select_m be the set of all such functions. We re-state (9) as follows:

For each moment m and each selection function s in Select_m ,

$$\bigcap_{\alpha \in \text{AGENT}} s(\alpha) \neq \phi.$$

Belnap comments (Belnap et al. 2001, p. 218) that while *Independence* is a “fierce” constraint (implying, for example, that no two agents can have the same possible choices at the same moment), it is also fairly weak (“banal” is Belnap’s term): it would be strange indeed if without causal priority, one agent’s choices could constrain what the other agent may choose.

The other important idea is *joint agency*. Again, the basic idea can be explained with just two agents. Let $\Gamma = \{\alpha, \beta\}$, where α and β are distinct agents. We want to define truth conditions for $[\Gamma \text{ cstit}: A]$. The concept is illustrated by referring once again to Fig. 3. With the assignments given by our model M , neither α nor β can see to it that A on any history in m . However, $M, m/h_1 \models [\Gamma \text{ stit}: A]$ because A holds at every history h' that belongs to the choice cell containing h_1 that is determined *jointly* by α and β . The formal definition is the same as (8) except that the condition invokes equivalence within the choice cell determined jointly by the agents in Γ .

3 Causal Independence

Both causal decision theory and Horty’s deontic logic depend upon the concept of causal independence. Before examining how causal independence can be characterized in the *stit* framework, it is helpful to review its role in causal decision theory. I focus on the formulation due to Skyrms (1980).⁶

Skyrms’s formulation requires the identification of a set of independent causal factors that provide the background for an agent’s choices. Each independent causal factor is represented as a random variable X_i . For simplicity’s sake, suppose that the set of possible outcomes O_1, \dots, O_n of interest to the agent is finite, that the set X_1, \dots, X_N of independent factors relevant to these possible outcomes is also finite, and that each variable X_i can take on finitely many values. Then the set S consisting of all possible combinations of assignments to these variables is also finite. This set constitutes a partition of the set of possible worlds into *causal background contexts* or *states* S_1, \dots, S_M : each state S_i is obtained by specifying one possible value for each of X_1, \dots, X_N . Suppose that the agent has a finite set $\{K_1, \dots, K_m\}$ of available alternative acts. The crucial idea is that the causal factors *in conjunction* with these alternative acts determine relevant conditional chances of the possible outcomes: the

⁶ There are numerous formulations of causal decision theory, including (Gibbard and Harper 1978), (Skyrms 1980) and (Joyce 1999). (Skyrms 1980) is in some ways the simplest and most relevant to our present concerns.

conditional chances $P(O_k / K_i \wedge S_j)$ are *constant* within each $K_i \wedge S_j$. Finally, the agent assigns a utility $u(O_k \wedge K_i \wedge S_j)$ to each outcome-act-state combination.

For a standard example, let K_i be the selection of a ball from one of m urns, let O_1, \dots, O_n represent n different colours of ball that may be drawn, and let S_1, \dots, S_M stand for M possible initial arrangements of coloured balls in the m urns. Then the probability $P(O_k / K_i \wedge S_j)$ is the conditional chance of drawing a ball of colour O_k , given arrangement S_j and the act K_i of drawing from urn i . The utility $u(O_k \wedge K_i \wedge S_j)$ depends upon the desirability of each combination (perhaps the agent has placed a bet in advance).

To allow for cases in which the agent is uncertain about the background context, Skyrms introduces a subjective probability distribution $prob(S_j)$ over all of the states. In the urn example, this represents your initial credence about the likelihoods of the different possible arrangements. The expected utility of act K_i is then given by the formula

(10) *Expected utility.*

$$\mathcal{U}(K_i) = \sum_j prob(S_j) \sum_k P(O_k / K_i \wedge S_j) \cdot u(O_k \wedge K_i \wedge S_j).$$

The thesis of causal decision theory is that *a rational agent maximizes expected utility* as given by (10). The equation highlights the importance of independent causal factors in the theory; the outer summation is over all possible states.⁷

Horty's deontic logic has a similar, but much weaker, guiding principle. Without conditional chances or credences, his analysis (outlined in the next section) is based solely upon the concept of dominance. Yet dominance reasoning shares with causal decision theory the need for a set of independent causal factors and a corresponding set of causal background contexts. *Stit* frames help to make these things precise by providing a concrete interpretation of possible worlds and a plausible way to identify some of the independent causal factors. I first review Horty's account and then propose a slight generalization.

Horty begins with an informal explication of causal independence:

...the basic intuition ... is that a proposition is supposed to be causally independent of the actions available to a particular agent whenever its truth or falsity is guaranteed by a source of causality other than the actions of that agent. (2001, 82)

Recall postulate (9), *Weak Independence of Agents*: if agents in a group make simultaneous choices, then the intersection of the relevant choice sets is non-empty. Horty strengthens this in two ways. First, he assumes that all choices at m by agents other than α are independent causal factors for α 's choice at m . This strengthened assumption, stated in counterfactual terms, applies to all moment-history pairs m/h and all agents α .

⁷ In this chapter, for the sake of simplicity, we ignore the element of subjective probability represented by *prob*. Skyrms (1994) provides a good discussion. In the present framework, subjective probability could usefully be introduced to represent the agent's uncertainty about location, i.e., about which m is the moment of decision.

(11) *Strong Independence (of Agents)*.

Let S represent the intersection of all actual choices (i.e., choice cells) of all agents other than α at m/h . If α were to make a different choice than the one made at m/h , the other agents would still (collectively) choose S .

Second, Horty adopts a provisional simplifying assumption that I shall refer to as *Causal Completeness of AGENT*.

(12) *Causal Completeness (of AGENT)*.

Choices by agents in $AGENT \setminus \{\alpha\}$ (i.e., agents other than α) are the *only* independent causal factors relevant to α 's choice.

Taken together, the two assumptions imply that the independent causal factors for α 's choice are precisely choices by other agents.

To illustrate these ideas, consider Fig. 3 again. Suppose that the picture represents choices by two agents, α and β , to cooperate in moving a heavy box at moment m . At m/h_1 and m/h_2 , the box is successfully moved (represented by the proposition A). Here, α chooses a_1 and β chooses b_1 . The choice by β is causally independent of the choice by α (by *Strong Independence*): if α were to choose a_2 (don't cooperate) at m , then β would still choose b_1 . Further, this is the only relevant independent causal factor for α 's choice (by *Causal Completeness*).


What defense can we give for assumptions (11) and (12)? For the first, the argument is that from *Weak Independence of Agents* and the simultaneity of choices by other agents, it is reasonable to infer *Strong Independence*. Simultaneous choices by agents must be causally independent of each other and of α 's choices at m .⁸ By contrast, *Causal Completeness* is offered merely as a useful "initial approximation" for Horty's deontic logic. Horty explicitly identifies two sources of independent causal influence that are not reflected in his account: "nonagentive sources" (*Nature*) and later choices by agents other than α (Horty 2001, 89–95). While I agree with Horty that fully to incorporate these influences into the analysis would be a "substantial research task", I believe that important special cases can be accommodated without great difficulty.

Consider first the case of *Nature*. Figure 4 illustrates a version of *Example 1 (Gambler)*, described at the start of this chapter.⁹ At moment m , α has a choice of gambling (G) or not. The gamble costs \$5. A fair coin toss is to be performed at m whether or not α gambles. If the coin comes up *Heads* (H), α leaves with \$10, but on *Tails* (T) she gets nothing. If α declines the gamble, she keeps her \$5. The outcomes

⁸ Talk of simultaneity suggests that there might be some gain in clarity by moving to a framework of *branching space-time* (Belnap 1992) instead of *branching time*. The added complexity of branching space-time is unnecessary, however, since for present purposes simultaneity is adequately characterized in a branching-time framework in terms of condition (7), *No Choice between Undivided Histories*. That is, it suffices that there exists a moment m such that for each agent, histories within the relevant choice cells are undivided at m while histories belonging to different choice cells are divided at m .

⁹ The *Gambler* example is due to Horty (2001), who formulates and discusses a number of versions.

Fig. 4 Gambler (I)



are as shown, where W signifies a win, L a loss and N the *status quo* where α neither wins nor loses.

There are no other agents besides α in this picture.¹⁰ Yet the coin toss has the characteristics of an independent causal factor. It occurs simultaneously with α 's choice. It satisfies an analogue of *Weak Independence*: any choice by α is compatible with either result, *Heads* or *Tails*. Finally, it is reasonable to regard the coin toss as satisfying *Strong Independence*. Consider

(13) If α had gambled, he would have won.

We endorse the truth of (13) at m/h_3 and m/h_4 , and its falsity at m/h_7 and m/h_8 . To summarize these observations, we introduce a random variable *Toss* with values $\{\text{Heads}, \text{Tails}\}$. With respect to both *Weak Independence* and *Strong Independence*, *Toss* has characteristics analogous to those of an agent with choice set $\{\text{Heads}, \text{Tails}\}$.

The generalization proposed here is to extend Horty's account of causal independence to include not just agents but also *chance mechanisms operating simultaneously with α 's choice*. By *chance mechanisms*, I mean well-understood processes such as those employed in games of chance: coin tosses, dice rolls, card drawings and so forth. These are singled out for two reasons. First, such processes have outcomes with well-defined and unproblematic probabilities.¹¹ Second, these processes are crucial in defining *mixed strategies*, which will be important later in this chapter. Each such mechanism can be modeled as a random variable X that may take different possible values $X = x_i$ at the moment m , i.e., at distinct moment history pairs m/h and m/h' . Let VAR be the set of independent variables representing chance processes.¹² We shall make assumptions about VAR that are entirely parallel to those for *AGENT*.

¹⁰ It may be that agent β tosses the coin, but β does not choose the result *Heads* or *Tails*. So the clusters shown are not choice sets for β .

¹¹ Gillies (2000) argues that such processes have a distinguished role in accounts of objective chance. In particular, the problem of identifying an appropriate reference class is relatively insignificant.

¹² We could relativize to each moment, using VAR_m for the set of variables that represent chance processes operating at m . We avoid this relativization both because we shall only ever be concerned

First, each random variable X must satisfy an analogue of the *No Choice Between Undivided Histories* condition, representing the fact that the chance process operates at moment m (rather than at a later moment). A pair of definitions makes this clear.

(14) $Rng^m(X)$.

By analogy with $Choice_\alpha^m$, if X is a random variable, define $rng^m(X)$ as the partition of H_m corresponding to the possible values $X = x_i$ at m . (Two histories h_1 and h_2 belong to the same element of $rng^m(X)$ if for some x_i , both $M, m/h_1 \models X = x_i$ and $M, m/h_2 \models X = x_i$.) For simplicity, we shall assume that these values x_i are always real numbers.¹³

(15) *No Separation of Undivided Histories*.

Whenever h_1 and h_2 are undivided at m , X has the same value $X = x_i$ at both m/h_1 and m/h_2 . That is, h_1 and h_2 must belong to the same element of $rng^m(X)$.

Condition (15) rules out random variables that partition H_m based on future processes.

Next, we need analogues for (9) *Weak Independence* and (10) *Strong Independence*. We want these analogues to apply to agents and chance processes taken together, which motivates the following definitions.

(16) *FACTOR*.

FACTOR is the union of *VAR* and the set of random variables representing choices by agents:

$$FACTOR = AGENT \cup VAR$$

(17) *Extended Selection Function*.

An *extended selection function* s at moment m is a mapping from *FACTOR* into H_m that selects a choice in $Choice_\alpha^m$ for each agent α in *AGENT*, and an element of $rng^m(X)$ for each variable X in *VAR*.

As before, we use the notation $Select_m$ for the set of all such functions.

It is sometimes convenient to regard the agents in *AGENT* as random variables, and to represent $Choice_\alpha^m$ as $rng^m(\alpha)$. This allows us to state a compact analogue of (9):

(Footnote 12 continued)

with a single moment and to maintain the analogy with *AGENT*, the set of agents fixed over all moments.

¹³ For the purposes of this chapter and to maintain consistency with the definitions in Sect. 2, we assume that all statements $X = x_i$ can be represented in our language as propositional constants. My thanks to Thomas Müller for pointing this out.

(18) *Weak Independence of FACTOR*.

For each moment m and each extended selection function s in $Select_m$,

$$\bigcap_{X \in FACTOR} s(X) \neq \phi.$$

We also have an obvious formulation of *Strong Independence*:

(19) *Strong Independence of FACTOR*.

Let S represent the background state for $X = x_i$ at m/h , i.e.,

$$S = \bigcap_{Y \in FACTOR \setminus \{X\}} s(Y),$$

where $s(Y)$ is the element of $rng^m(Y)$ selected for Y at m/h . If a different value $X = x_j$ were selected at m , the other variable values and hence the background state S would remain the same.

[In particular, if any agent α were to make a different choice, all choices by other agents and all variable values would remain the same].¹⁴

On this account, causal independence extends to chance mechanisms that operate independently of each other and of agents, such as the coin toss in *Gambler* (I). We acknowledge this modification by extending our earlier definition of *stit frames*.

(20) *Extended stit frames*.


An *extended stit frame* is a structure $\langle Tree, <, FACTOR, Rng \rangle$ that satisfies all earlier assumptions as well as (15), (18) and (19).

We are not quite done. A separate approach is needed to represent chance mechanisms *initiated* by agents. Consider a variation of *Gambler*: α tosses the coin if and only if she accepts the gamble; if she declines, there is no coin toss (Fig. 5).

In this case, it is inappropriate to model *Toss* as an independent causal factor. There is no independent partition of H_m ; the toss does not happen if α declines to gamble. So both *Weak Independence* and *Strong Independence* fail.

The difficulty is that while *Gambler* (II) involves a well-understood chance mechanism that (to paraphrase Horty) represents a source of causality other than the actions of α , the mechanism does not operate independently of α and cannot be modelled as a random variable in *VAR*. An alternative approach, following Skyrms, is to represent

¹⁴ Note that the plausibility of (19) depends upon the modest scope of the set *VAR*. The stated (though still undeniably vague) restriction is that the random variables in *VAR* are restricted to well-understood chance mechanisms, the sort of mechanisms that one could exploit in implementing a mixed strategy (see Sect. 7). In particular, I mean to exclude quantum phenomena.

Fig. 5 Gambler (II)

the operation of such mechanisms via conditional chances for *outcomes* within each causal background context. This approach will be developed below in Sect. 6.

4 Horty's Dominance Ought

Consider *Gambler* (I), as illustrated in Fig. 4. Substitute numerical values 10 in place of *W* (a winning gamble), 0 in place of *L* (a loss), and 5 in place of *N* (no gamble). These values represent the money that agent α possesses when the dust settles. We can also think of them as utilities that represent α 's preferences.

Adding utilities to a *stit frame* gives us a *utilitarian stit frame*, defined by Horty as a structure of the form

$$\langle Tree, <, AGENT, Choice, Value \rangle,$$

where *Tree*, *<*, *AGENT* and *Choice* are as in Sect. 2, and *Value* is a function that assigns a real number $Value(h)$ to each history. A *utilitarian stit model* combines a utilitarian stit frame with an assignment of truth values to propositions. Figure 6 illustrates *Gambler* (I) in a utilitarian stit model.


Horty provides a semantics for statements of the form

$$\odot[\alpha \text{ cstit}: A] \quad (\alpha \text{ ought to see to it that } A).^{15}$$

The basic idea of his “dominance ought” is that α ought to see to it that *A* iff *A* is guaranteed by *each* optimal (non-dominated) choice. It takes care to make this

¹⁵ Horty uses \odot to distinguish his *dominance ought* from other obligation operators. I shall use \preceq and \prec to represent the corresponding dominance relations, described below. This helps to distinguish Horty's dominance ordering from variants to be introduced in later sections.

Fig. 6 Gambler (I) (utilitarian stit model)



precise and to handle cases where there are no optimal choices. The elements of Horty's account include *background states*, a *value ordering* on propositions at m (subsets of H_m), and the *dominance* relation between possible choices for an agent.

(21) *Dominance ordering on choices.*

- $State_\alpha^m$: the partition of histories through m into background causal contexts S for α 's choice. For Horty, as we have seen, these background contexts are simply joint choices by all members of *AGENT* other than α .
- *Ordering relations* (\leq and $<$) on propositions at moment m : If X and Y are two subsets of H_m , then (1) $X \leq Y$ if $Value(h) \leq Value(h')$ for each $h \in X$ and $h' \in Y$, and (2) $X < Y$ if $X \leq Y$ and in addition, $Value(h) < Value(h')$ for some $h \in X$ and $h' \in Y$.
- *Dominance relations* (\preceq and \prec) on $Choice_\alpha^m$: If K and K' are members of $Choice_\alpha^m$ (i.e., possible choices for α at m), then (1) $K \preceq K'$ (K' weakly dominates K) if $K \cap S \leq K' \cap S$ for each state S in $State_\alpha^m$, and (2) $K \prec K'$ (K' strongly dominates K) if $K \preceq K'$ and, in addition, $K \cap S < K' \cap S$ for some state S in $State_\alpha^m$.
- *Optimal acts*. If $K \in Choice_\alpha^m$ is a possible act for α at m , and there is no $K' \in Choice_\alpha^m$ such that $K \prec K'$, then K is an *optimal act* for α at m .

To illustrate these ideas, imagine that in Fig. 6, the result *Heads* or *Tails* is determined by another agent, β . In this case, $State_\alpha^m$ is $\{Heads, Tails\}$, and $(G \& Tails) < \sim G < (G \& Heads)$ on the propositional ordering. Neither $[\alpha \text{ cstit: } G]$ nor $[\alpha \text{ cstit: } \sim G]$ is a dominated act for α , so both are optimal.

In simple cases where α has only finitely many possible choices, Horty's account of obligation is as follows:

(22) *Horty Dominance Ought (Finite Choice case).*

$M, m/h \models \odot[\alpha \text{ cstit: } A]$ iff $M, m/h' \models A$ for all h' belonging to any choice K that is optimal at m .

That is, α ought to see to it that A iff every optimal choice *guarantees* A . In the case of *Gambler*, the Horty account tells us that neither $\odot[\alpha \text{ cstit}: G]$ nor $\odot[\alpha \text{ cstit}: \sim G]$ is true. In the absence of probabilistic information, gambling and not gambling are both permitted.

There may be situations where α has infinitely many options, none of which is optimal. For instance, if α and β are playing the *greatest integer game*, where the person who names the largest integer wins, then there is no optimal choice. In such cases, there are still *dominated* choices and hence there are still obligations—for instance, the obligation to choose an integer greater than 1,000. To accommodate such cases, Horty provides a more general evaluation rule.

(23) *Horty Dominance Ought (general case).*

$M, m/h \models \odot[\alpha \text{ cstit}: A]$ iff for each choice $K \in \text{Choice}_\alpha^m$ that does not guarantee A , there is a choice $K' \in \text{Choice}_\alpha^m$ such that (1) $K \prec \cdot K'$ (K' strongly dominates K), (2) $M, m/h' \models A$ for all h' belonging to K' , and (3) for every choice $K'' \in \text{Choice}_\alpha^m$ such that $K' \preceq \cdot K''$, $M, m/h'' \models A$ for all h'' belonging to K'' .

The requirement for $\odot[\alpha \text{ cstit}: A]$ is that any action K that does not guarantee A is dominated by an action K' that does guarantee A and is either optimal or dominated only by other actions that guarantee A .

Suppose we modify *Gambler* so that the values are as shown in Fig. 7. Once again, we imagine that the result *Heads* or *Tails* is determined by the choice of another agent, β , so that the background states for α 's choice are $\{\text{Heads}, \text{Tails}\}$.


As before, gambling is not always better than not gambling: the value of h_3 and h_4 exceeds that of h_5 and h_6 . But this time, the act of gambling dominates the act of not gambling, so that $\odot[\alpha \text{ cstit}: G]$ is true at m .¹⁶

As a slight modification of Horty's account, let us bring in the generalization of causal independence introduced in the *extended stit frames* of Sect. 3. Suppose that, in Fig. 7, the result *Heads* or *Tails* is determined by a chance mechanism (a coin toss) rather than by another agent's choice. Since Horty's causal background contexts only take into account choices by other agents, his account gives us no partition between *Heads* and *Tails*, no dominance of $[\alpha \text{ cstit}: G]$ over $[\alpha \text{ cstit}: \sim G]$, and hence no obligation to gamble. The extension of *AGENT* to *FACTOR*, as explained in Sect. 3, remedies this problem by counting chance mechanisms such as coin tosses as independent causal factors on par with choices. If we allow this extension, then $\text{State}_\alpha^m = \{\text{Heads}, \text{Tails}\}$ in Fig. 7, which restores the dominance reasoning that leads to $\odot[\alpha \text{ cstit}: G]$. Henceforth, $\preceq \cdot$ and \odot will be understood as incorporating this extended concept of factors and background states. Other than the change to State_α^m , there is no formal modification required for definitions (21), (22) and (23).¹⁷

¹⁶ Technically, true at all m, h pairs. Since the semantics guarantees that $\odot[\alpha \text{ cstit}: G]$ is either settled true or settled false at a moment, however, we may speak of obligations as holding at a *moment*.

¹⁷ We could similarly define *extended utilitarian stit frames* by substituting *FACTOR* for *AGENT* and *Rng* for *Choice*.

Fig. 7 Gambler III (utilitarian stit frame)



The distinctive feature of Horty's approach, in contrast to a great deal of earlier work in deontic logic, is that his semantics for obligation is based on an ordering on *choices*, rather than an ordering on *histories* or *worlds*. Horty's deontic logic gives us a weak decision theory, namely, the part of decision theory that corresponds to dominance reasoning. Let us say that an ordering \preccurlyeq on choices K in $Choice_\alpha^m$ is *admissible* if it extends Horty's dominance ordering: $K \preccurlyeq K'$ whenever $K \prec K'$. Any decision principle based upon an admissible ordering on choices will preserve obligations that hold according to the *dominance ought*. But the reverse is not true: stronger decision principles justify assertions of obligation that fail on the Horty semantics. The remainder of this chapter shows how three of these stronger decision principles, and the corresponding notions of obligation, can be modeled by extensions within Horty's framework.

5 Decisions Under Ignorance: The Maximin Ought

In this section, I show how Horty's account might be extended to incorporate a principle that is sometimes used for making decisions under ignorance: *maximin*. The maximin rule tells the agent to compare *minimum* utilities possible for each available act, and to choose the act with the maximal minimum utility. The rationale behind this rule is *conservatism*: by following maximin, the agent guarantees the least bad outcome. In *Gambler (I)*, for instance, the choice "Don't Gamble" guarantees a utility of 5, while "Gamble" allows possible outcomes with utilities of 0 and 10. Maximin thus prescribes the choice of not gambling. By contrast, as we have just seen, Horty's dominance ought prescribes nothing, since both gambling and not gambling are optimal choices.

There is an important ambiguity in the phrase “decisions under ignorance.” Commonly, such decisions are characterized as those made “when it makes no sense to assign probabilities to the outcomes emanating from one or more of the acts” (Resnik 1987, 14). We can distinguish between cases of *total ignorance*, where the agent has no probabilistic information at all (not even knowledge of independence), and *ignorance of probabilities*, where the agent has no quantitative probabilistic information but does possess knowledge of causal independence. The latter case, in which the agent has exactly the same information as required for the dominance ought, is our focus. The objective is to provide a semantics for *maximin ought-to-do*, O_m , that strengthens Horty’s dominance ought in the following sense:

$$(24) \odot[\alpha \text{ cstit}: A] \models O_m[\alpha \text{ cstit}: A].$$

Both operators are formulated within utilitarian stit frameworks. The meaning of (24) is that for any utilitarian stit model M and for any m, h pair, if $M, m/h \models \odot[\alpha \text{ cstit}: A]$ then $M, m/h \models O_m[\alpha \text{ cstit}: A]$.

Unfortunately, our preliminary statement of the maximin rule is inconsistent with Horty’s dominance ought. To see this, consider a kid-friendly version of *Gambler* that rewards a decision to gamble with \$10 if *Heads*, \$5 if *Tails*; a decision not to gamble yields \$5 regardless of outcome.¹⁸ Gambling is plainly the dominant act. Yet the simple maximin rule regards gambling and not gambling as equally good because the worst outcome on either choice is \$5. This violation of dominance can be avoided by moving to a lexical version of maximin,¹⁹ but an alternative approach will be offered below.

Another weakness of maximin as stated is its inability to handle a situation of infinite choices, such as the *greatest integer game* discussed in the preceding section. Even though no available act attains a maximal minimum value, it seems clear that maximin reasoning should license many of the same conclusions as dominance reasoning—for instance, that one ought to select an integer greater than 1,000.

We proceed in stages, starting with a new ordering on choices that *combines* maximin with the dominance ordering $\preceq \cdot$ defined in (21). The idea is to apply maximin only to pairwise comparisons where neither choice dominates the other.

(25) *Non-dominance*. If K and K' are members of Choice_α^m (i.e., possible choices for α at m), write $K \not\preceq K'$ if neither $K \preceq K'$ nor $K' \preceq K$.

(26) *Maximin ordering* (\preceq_m and $<_m$) on Choice_α^m :

If K and K' are members of Choice_α^m (i.e., possible choices for α at m), then

- (1) $K \preceq_m K'$ if (i) $K \preceq K'$ or (ii) $K \not\preceq K'$ and $\inf\{\text{Val}(h)/h \in K\} \leq \inf\{\text{Val}(h')/h' \in K'\}$; and

¹⁸ Kid-friendly because the gambler never loses any money.

¹⁹ See (Resnik 1987).

- (2) $K \prec_m K'$ if (i) $K \prec \cdot K'$ or (ii) $K \not\leq K'$ and $\inf\{Val(h)/h \in K\} < \inf\{Val(h')/h' \in K'\}$.²⁰

(27) *Maximin ought*, O_m .

$M, m/h \models O_m[\alpha \text{ cstit}: A]$ iff for each choice $K \in \text{Choice}_\alpha^m$ that does not guarantee A , there is a choice $K' \in \text{Choice}_\alpha^m$ such that (1) $K \prec_m K'$, (2) $M, m/h' \models A$ for all h' belonging to K' , and (3) for every choice $K'' \in \text{Choice}_\alpha^m$ such that $K' \preccurlyeq_m K''$, $M, m/h'' \models A$ for all h'' belonging to K'' .

The relationship (24), that $\odot[\alpha \text{ cstit}: A]$ entails $O_m[\alpha \text{ cstit}: A]$, is clear because dominance is built into the definition of O_m . By way of example: in the kid-friendly version of *Gambler*, there is an obligation to gamble because gambling dominates not gambling. In the original version of *Gambler* (illustrated in Fig. 6) there is no dominant choice, but not gambling is superior to gambling on the maximin ordering; as a consequence, the agent has an obligation not to gamble ($O_m[\alpha \text{ cstit}: \sim G]$). The same result holds when the result of *Heads* or *Tails* is achieved through placement of the coin by an independent agent. Finally, consider an infinite choice situation such as the *Greatest Integer Game*, where each possible choice of an integer is dominated by any choice of a larger integer. By (27), it is still true that one ought to choose an integer larger than 1,000.

The formulation of *maximin ought* in (27) has some advantages over the traditional formulation of *maximin* in decision theory. The first is its compatibility with dominance reasoning. The standard version of maximin, as noted earlier, does not always exclude dominated choices; the same problem applies to some forms of *lexical maximin* reasoning.²¹ Other versions of lexical maximin, which respect dominance, are defined only for finite choice situations. By contrast, (27) is defined for arbitrary choice situations and is always compatible with dominance reasoning. A second advantage of the present formulation, indeed, is its ability to accommodate infinite choice situations, as noted in the preceding paragraph. In infinite choice situations where no individual choice is rational, we can still identify obligations. This highlights a general advantage of locating decision principles within deontic logic: whereas decision theory is focused specifically on rational *acts*, deontic logic provides truth conditions for all sentences of the form $O_m[\alpha \text{ cstit}: A]$.

The point of this discussion is not to endorse the *maximin ought* over Harty's *dominance ought*. The weaknesses of maximin reasoning are well known.²² There are two motives for developing O_m . The first is simply to flesh out the claim that the Harty semantics can be strengthened to yield a stronger decision theory. The second is that maximin reasoning plays an important role in game theory (Sect. 7).

²⁰ If S is a set of real numbers that is bounded below, then $\inf(S)$ refers to the *infimum* or greatest lower bound of S . Thus, I assume that the set of utility values within each K is bounded. The assumption of bounded utilities is standard in decision theory, to avoid problems such as the St. Petersburg paradox (see Resnik 1987, p. 107). Here, we require only the weaker assumption that utilities are bounded below within each possible choice.

²¹ See (Resnik 1987).

²² See (Resnik 1987) for discussion.

6 Decision Under Risk: Probabilistic Utilitarian stit Frames

This section extends Horty's account in a different direction by incorporating a simple type of probabilistic information – that which is related to chance mechanisms — into the semantics of obligation. This results in a strengthening of the dominance ought that is incompatible with maximin (just as expected utility reasoning is incompatible with maximin reasoning in decision theory). For simplicity, we ignore other agents; we have a single agent, α , making choices. We assume that α has finitely many choices and that there are only finitely many relevant independent causal factors, so both $Choice_\alpha^m$ and $State_\alpha^m$ are finite.

Let us begin with *Gambler* (I) as depicted in Figs. 4 (without utilities) and 6 (with utilities). Suppose that we have a coin toss with known probabilities 0.5 for *Heads* and *Tails*. In the theory of decision under risk, a straightforward application of expected utility reasoning yields a tie: gambling and not gambling have equal expected utility ($EU = 5$). A similar analysis yields a tie for *Gambler* (II) as depicted in Fig. 5, where the coin toss only occurs if α decides to gamble. But it is clear in these examples that slight changes to the utilities would tip the decision one way or the other. To extend Horty's account to such cases, we need to add some concepts to utilitarian *stit* frames. It suffices to add two additional concepts: *outcomes* and *conditional chances*.

We shall assume a finite set O_1, \dots, O_n of *outcomes* of interest. These are propositions at moment m (subsets of H_m) that constitute a partition of H_m and which, in conjunction with the background contexts and the agent's choices, influence the assignment of conditional chance and utility. In particular, they allow us to represent probabilistic information about chance processes initiated by agents; such processes cannot be treated as independent causal factors, as explained at the end of Sect. 3. In *Gambler*, the outcomes may be described as $\{Win, Lose, Neither\}$.

For the *conditional chance* function on H_m , the simplest approach is to take the underlying algebra²³ of subsets of H_m to consist of all finite unions of sets $K_i \wedge S_j \wedge O_k$, where K_1, \dots, K_m are available acts, S_1, \dots, S_M are the background contexts, and O_1, \dots, O_n are the outcomes. Probabilistic information about chance mechanisms is given by a conditional chance function P , assigning values $P(O_k/K_i \wedge S_j)$ that we take as primitive. P must satisfy the standard axioms of the probability calculus. Since the algebra is finite, P need only be finitely additive.²⁴

The following two assumptions would allow easy extension of Horty's framework to handle probabilistic choices:

- (a) *Uniform conditional chances.* $P(O_k/K_i \wedge S_j)$ is constant for each relevant outcome O_1, \dots, O_n within each $K_i \wedge S_j$.
- (b) *Uniform utilities.* $Val(h) = Val(h')$ for all $h, h' \in K_i \wedge S_j \wedge O_k$, for all i, j, k .

²³ An algebra of subsets of X is a family \mathcal{F} of subsets that includes X and the empty set, and is closed under finite unions, intersections and complementation.

²⁴ The function P may, but need not, assign unconditional chances to elements of the algebra, including α 's own actions. Since P represents objective chance, difficulties alleged to exist for the assignment of subjective probabilities to one's current choices are not relevant; see (Levi 1997) and (Spohn 1977).

Given these assumptions, we can “import” decision theory into the *stit* framework. We can use expected utility maximization as the criterion for what agent α ought to do at m , in simple cases such as *Gambler* (I) and (II).

But these assumptions need not always hold. First, assumption (a) might fail. Some of the objective chances used in decision theory are not conditional chances associated with chance mechanisms. For example, they may be derive from observed frequencies. So there may be information about objective conditional chances that is not represented in the *stit* framework. Second, assumption (b) might fail. Within a single choice-state combination, we might find histories with different utilities based (for example) on future choices by agents or future events. In general: if it is impossible to find a set of outcomes satisfying (a) and (b), then it is impossible to apply straightforward expected utility maximization. To keep matters simple, however, I shall assume that condition (a) is satisfied but that (b) may fail.

This leads us to a definition of *probabilistic utilitarian stit frames*.

(28) A *probabilistic utilitarian stit frame* is a structure of the form.

$$\langle Tree, <, FACTOR, Rng, Value, Outcome, P \rangle,$$


where *Tree*, *<*, *FACTOR*, *Rng* and *Value* are as in Sects. 2, 3, and 4, *Outcome* is a function that assigns to each moment m a partition $\{O_1, \dots, O_n\}$ of H_m and $P(\cdot/\cdot)$ is a conditional probability function that assigns a value $P(O_k/K_i \wedge S_j)$ for each outcome O_k , choice K_i and state S_j .²⁵

A *probabilistic utilitarian stit model* combines a probabilistic utilitarian stit frame with an assignment of truth values to propositions.

In order to formulate the concept of obligation in probabilistic utilitarian stit frames, consider the following case (Fig. 8). In this example, the coin toss is an independent factor and, as usual, *Heads* and *Tails* have fixed conditional chances of 0.5 regardless of whether α gambles. The outcomes are *Win*, *Lose* and *Neither*, but this time the utilities of *Win* and *Lose* are not fixed (i.e., assumption (b) fails). So there is no sharp value for the expected utility of gambling. Still, we can see that the expected utility of gambling is at least $(0.5)(9) + (0.5)(2) = 5.5$, which exceeds the expected utility of not gambling. Thus, we ought to gamble.

This motivates the following account of obligation, replacing dominance with *dominating expectation* in the Horty semantics. We continue to assume that both $Choice_\alpha^m$ and $State_\alpha^m$ are finite.

²⁵ One other assumption is necessary: the utilities given by *Value* (or *Val*) represent an interval scale (i.e., they are unique up to a positive linear transformation). This assumption guarantees that the expected utility ordering on choices, defined below, is invariant under allowable changes in representation of the utilities. Consider Fig. 8 below, which depicts utilities assigned by a particular function *Val*. The expected utility calculation given below, which shows that we ought to gamble, fails if the agent’s utilities are equally well represented by a function *Val'* that assigns 7 to h_1 , 6 to h_2 , and keeps all other values the same as *Val*. Although *Val* and *Val'* agree on their ordinal ranking of histories, they are not related by a positive linear transformation. If $Val' = aVal + b$ for $a > 0$, however, then they induce the same ordering on choices.

Fig. 8 Gambler (IV)

(29) *Dominating expectation ordering* (\preceq_d and \prec_d) on $Choice_\alpha^m$.

If K and K' are members of $Choice_\alpha^m$ (i.e., possible choices for α at m), then

- (1) $K \preceq_d K'$ if for all choices h_{jk}, h_{jk}' of histories in $K \cap S_j \cap O_k$ and $K' \cap S_j \cap O_k$ respectively, $\Sigma_{j,k} Val(h_{jk})P(O_k/K \wedge S_j) \leq \Sigma_{j,k} Val(h_{jk}')P(O_k/K' \wedge S_j)$; and
- (2) $K \prec_d K'$ if $K \preceq_d K'$ but not $K' \preceq_d K$.

This principle says that act K' is better than act K if the expectation of K' dominates the expectation of K .

Corresponding to this new ordering on choices, we have a new concept of obligation defined analogously to the earlier definitions (23) and (27). It states, roughly, that $O_d[\alpha \text{ cstit}: A]$ if A is guaranteed by all choices whose expectation is not dominated.

(30) *Dominating expectation ought*, O_d .


$M, m/h \models O_d[\alpha \text{ cstit}: A]$ iff for each choice $K \in Choice_\alpha^m$ that does not guarantee A , there is a choice $K' \in Choice_\alpha^m$ such that (1) $K \prec_d K'$,

- (2) $M, m/h' \models A$ for all h' belonging to K' , and (3) for every choice $K'' \in Choice_\alpha^m$ such that $K' \preceq_d K''$, $M, m/h'' \models A$ for all h'' belonging to K'' .

It should be clear that O_d is a strengthening of Horty's \odot , since $K \preceq_d K'$ whenever $K \preceq \cdot K'$.²⁶ On the other hand, just as we would expect, O_d diverges from the *maximin ought* O_m . In the version of *Gambler* depicted by Fig. 8, we have $O_d[\alpha \text{ cstit}: G]$ while $O_m[\alpha \text{ cstit}: \sim G]$.

²⁶ Proof: if $K \preceq \cdot K'$, then for any state S_j , $Val(h) \leq Val(h')$ for each $h \in K \cap S_j$ and $h' \in K' \cap S_j$. Then $\Sigma_k Val(h_{jk})P(O_k/K \wedge S_j) \leq \Sigma_k Val(h_{jk}')P(O_k/K' \wedge S_j)$ for all choices $h_{jk} \in K \cap S_j \cap O_k$ and $h_{jk}' \in K' \cap S_j \cap O_k$, and the result follows.

Fig. 9 Matching pennies (a) stit picture; (b) game theory



7 Game Theory and Mixed Strategies

Finally, I consider briefly how Horty's account might be extended to handle obligations in the setting of game theory. To do this in general would introduce many complications, including the need for separate *Value* functions to keep track of each agent's utilities.²⁷ My interest here lies mainly in showing how we might use the probabilistic utilitarian stit frames of the preceding section to make sense of *mixed strategies* in game theory. For this reason, I limit the discussion to *two-person zero-sum games*. The *Value* function represents α 's utilities, while utilities for the other agent, β , are exactly the negative of α 's utilities. We initially assume that finitely many choices—*pure strategies* in game theory—are available to both agents, and that there are no additional independent causal factors. Thus, the background contexts in $State_\alpha^m$ are simply β 's possible choices.

By way of motivation, notice that a very simple game, *Matching Pennies* (introduced at the outset of this chapter), generates difficulties for Horty's account. In this game, both α and β simultaneously display a coin with one side up. If the two displayed sides match (both *Heads* or both *Tails*), then β pays \$1 to α ; if the sides do not match, then α pays \$1 to β . The situation is illustrated in Fig. 9, with the *stit* and game-theoretic representations side by side.

In this situation, neither choice by α is optimal. Consequently, on Horty's account, we then have neither $\odot[\alpha \text{ cstit} : \text{Heads}]$ nor $\odot[\alpha \text{ cstit} : \text{Tails}]$. That is reasonable if the only choices available are the pure strategies $[\text{Display}] \text{Heads}$ or $[\text{Display}] \text{Tails}$. But Horty's conclusion is not plausible if we allow mixed strategies of the form

Display Heads with probability p and *Display Tails* with probability $1 - p$, abbreviated as

$$[p \text{ Heads}, (1 - p) \text{ Tails}]$$

or more simply as

²⁷ See Kooi and Tamminga (2008) for an account developed along these lines.

B_1	4	2
B_2	5	6
	A_1	A_2

Fig. 10 A zero-sum game

p Heads.

In game theory, this type of problem is solved by finding a *Nash equilibrium*: a pair of choices such that neither player can do better by unilaterally changing his or her choice. In *Matching Pennies*, there is a unique Nash equilibrium where both agents adopt the mixed strategy: $1/2$ *Heads*. This is the unique rational choice on the assumption that each player has full knowledge of the game and adopts the best possible strategy. In the remainder of this section, we suggest one way in which Horty's account can be expanded to accommodate mixed strategies, and then propose a semantics that yields the obligation to adopt an equilibrium strategy.

But first let's consider a preliminary question. How well does Horty's account fare if we limit ourselves to two-person zero-sum games with only pure strategies? Consider the following game (Fig. 10), with utilities for α shown.

Here, α chooses between the left (A_1) and right (A_2) columns, while β chooses between the top (B_1) and bottom (B_2) rows. From α 's point of view, neither choice is dominant: A_1 does better if β chooses B_1 , while A_2 does better if β chooses B_2 . So Horty's "dominance ought" yields no obligation: neither $\odot[\alpha \text{ cstit}: A_1]$ nor $\odot[\alpha \text{ cstit}: A_2]$ is true. However, both players will recognize that the top row is the dominant choice for β (whose utilities are the negative of those in Fig. 10). Given that β will choose B_1 , α ought to choose A_1 (Indeed, A_1 and B_1 constitute the unique Nash equilibrium for this game). This example shows that even without mixed strategies, Horty's dominance ought is inadequate for game theory.

One promising possibility might be the *maximin ought* (O_m) of Sect. 5, which combines dominance with *maximin* reasoning. Applied to Fig. 10, O_m gives the correct result: A_1 guarantees the maximal minimum, so $O_m[\alpha \text{ cstit}: A_1]$. Further encouragement comes from a standard result of game theory (Resnik 1987, p. 130):

Minimax equilibrium test. In a two-person zero-sum game, a necessary and sufficient condition for a pair of (pure) strategies to be in (Nash) equilibrium is that the payoff determined by them equal the minimal value of its (column) and the maximal value of its (row).²⁸ The values for all such equilibrium pairs are the same.

It is easy to establish the following proposition:

²⁸ Most expositions of game theory represent the utilities of the *Row* player in zero-sum games. Following *stit* conventions, I represent instead the utilities of the column player. The statement of the *Minimax equilibrium test* in (Resnik 1987) thus reverses *row* and *column*.

B_1	2	1
B_2	0	5
	A_1	A_2

Fig. 11 Another zero-sum game

- (31) *Proposition:* If there is a pure-strategy equilibrium pair in a two-person zero-sum game and if A is true at all non-dominated choices K for α that belong to such an equilibrium pair, then $O_m[\alpha \text{ cstit}: A]$.

Proof:

For each such K , $K' \preceq_m K$ for all choices K' ; thus, these K are *optimal* with respect to the ordering \preceq_m and we have $O_m[\alpha \text{ cstit}: A]$.

(31) shows that whenever there is a pure Nash equilibrium, the *maximin ought* correctly prescribes choices that are part of such an equilibrium.

Despite this success, *maximin ought* appears to overshoot the mark. It prescribes choices even when there is no pure equilibrium. Consider the following example (Fig. 11). Here there are no dominant choices for either player and no equilibrium. Nevertheless, *maximin ought* prescribes A_2 for α and B_1 for β , since these choices maximize minimal utility.

It might be interesting to consider whether there is any merit to these prescriptions. It might also be worth investigating whether there is a notion of obligation, intermediate in strength between \odot and O_m , that corresponds precisely to acts that comprise a Nash equilibrium. I pass over such investigations for the following reason: once we allow *mixed strategies*, the problem of capturing Nash equilibria with a Horty-style account of obligation is solved through an interesting combination of *maximin* reasoning and the weak concept of expected utility introduced in Sect. 6.

The first task is to give an analysis of mixed strategies. In game theory, a mixed strategy is commonly characterized as the use of a chance mechanism to select a pure strategy, followed by acting on the selected strategy. The details may not matter much in game theory, but they matter a great deal in the *stit* framework. If the chance mechanism operates at a moment prior to the choice of the pure strategy, then the analysis of a mixed strategy will involve both the prior moment when the mechanism operates and alternative later moments at which the agent chooses a pure strategy. To make things worse, the *stit* picture for the later moment will be identical to the original ‘pure strategy’ picture. If we evaluate the obligation at that later moment, it is unclear how the earlier operation of a chance mechanism can make any difference.

Perhaps the simplest approach, and the one which will be adopted here, is to represent each available mixed strategy as a separate choice existing at the same moment as the pure strategies. It is the choice of a chance mechanism whose possible outcomes are identical in structure with the pure strategies. Strictly speaking, this

Fig. 12 Matching pennies with mixed strategies

β	Heads	1	-1	1	-1
	Tails	-1	1	-1	1
β	Heads	1	-1	1	-1
	Tails	-1	1	-1	1

α

analysis requires that we modify $Choice_{\alpha}^m$ by adding one additional choice for each available chance distribution over the (finitely many) pure strategies. In practice, it usually suffices to represent all of the pure strategies plus a single choice that stands for an arbitrary mixed strategy (incorporating probabilistic parameters) or, on occasion, for a particular mixed strategy. Figure 12 illustrates *Matching Pennies* with mixed strategies. Dotted lines are used to separate outcomes for the case of choices that involve chance mechanisms.

None of the concepts of obligation described above gives the correct result here, namely, the obligation to choose $1/2$ Heads. According to the *dominance ought*, there are no obligations. The same is true for the *maximin ought*, since each choice has the same worst case. The *dominating expectation ought* is not even defined for settings involving multiple agents.

A helpful way to obtain the right choice ordering and the right concept of obligation is to exploit a well-known result from game theory (Resnik 1987, p. 136):

Maximin Theorem for two-person zero-sum games.

For every two-person zero-sum game there is at least one strategy (mixed or pure) for Row and at least one strategy for Col that form an equilibrium pair. If there is more than one such pair, their expected utilities are equal.

The expected utility for the equilibrium pair is referred to as the *security level* for both players because, by playing the equilibrium strategy, each player maximizes his or her minimum expected utility. The security level in *Matching Pennies* is $1/2$, which can be guaranteed by playing $1/2$ Heads. In contrast to our earlier discussion of zero-sum games with pure strategies, the inclusion of mixed strategies ensures the existence of an equilibrium.

The right ordering, then, is that one mixed strategy is preferable to another if its minimal expected utility exceeds that of the other. This ordering can be defined

within the setting of *probabilistic utilitarian stit frames*. Write P_α and P_β for α 's and β 's choice of mixed strategy. P_α and P_β are probability distributions over the choices available to α and β , respectively. That is, if K_1, \dots, K_m are the available pure strategies for α , i.e., the members of Choice_α^m , then $P_\alpha(K_i) = p_i$ with $\sum p_i = 1$; similarly, $P_\beta(B_j) = q_j$ for pure strategies B_1, \dots, B_n . The choice of a pure strategy K_i or B_j is just the special case where $p_i = 1$ or $q_j = 1$. Let $^*\text{Choice}_\alpha^m$ be the set of mixed strategies P_α based on the pure strategies in Choice_α^m .

(32) *Equilibrium ordering* (\preceq_e and \prec_e) on $^*\text{Choice}_\alpha^m$.

If P_α and P'_α are members of $^*\text{Choice}_\alpha^m$ (i.e., mixed strategies for α at m), then

- (1) $P_\alpha \preceq_e P'_\alpha$ if
 $\inf\{\sum_{i,j} \text{Val}(h_{ij}) P_\alpha(K_i) P_\beta(B_j) / h_{ij} \in K_i \cap B_j \text{ and } P_\beta \text{ a mixed strategy for } \beta\} \leq \inf\{\sum_{i,j} \text{Val}(h_{ij}) P'_\alpha(K_i) P_\beta(B_j) / h_{ij} \in K_i \cap B_j \text{ and } P_\beta \text{ a mixed strategy for } \beta\}$;

and

- (2) $P_\alpha \prec_e P'_\alpha$ if $P_\alpha \preceq_e P'_\alpha$ but not $P'_\alpha \preceq_e P_\alpha$.

The mixed strategy P'_α is better than P_α if it has greater minimal expected utility. The ordering \preceq_e is admissible in the following special sense: if $P_\alpha \preceq P'_\alpha$ where both P_α and P'_α are *pure strategies*, then $P_\alpha \preceq_e P'_\alpha$.²⁹

(33) *Equilibrium ought*, O_e .

$M, m/h \models O_e[\alpha \text{ cstit}: A]$ iff for each $P_\alpha \in ^*\text{Choice}_\alpha^m$ that does not guarantee A , there is a $P'_\alpha \in ^*\text{Choice}_\alpha^m$ such that (1) $P_\alpha \prec_e P'_\alpha$, (2) $M, m/h' \models A$ for all h' belonging to P'_α , and (3) for every $P''_\alpha \in ^*\text{Choice}_\alpha^m$ such that $P'_\alpha \preceq_e P''_\alpha$, $M, m/h'' \models A$ for all h'' belonging to P''_α .

In two-person zero-sum games where an equilibrium exists, (33) states that $O_e[\alpha \text{ cstit}: A]$ if A is guaranteed by all equilibrium mixed strategies. O_e is a strengthening of Horty's \odot , and it gives the right answer in the case of *Matching Pennies*: $O_e[\alpha \text{ cstit}: P_{1/2}]$.

Summarizing: mixed strategies can be defined in Horty's framework, and we can give an ordering on mixed strategies that yields the correct account of what agents ought to do in two-person zero-sum games. Extending these ideas to games involving more than two agents and to cooperative games may or may not be feasible.

8 Conclusion

Horty observes that his account of obligation “closes the gap” between deontic logic and act utilitarianism. That gap existed so long as deontic logic was viewed as an

²⁹ Proof: $P_\alpha(K_i) = 1$ for some i , and $P'_{\alpha'}(K'_{i'}) = 1$ for some i' . If $P_\alpha \preceq P'_{\alpha'}$, then $\text{Val}(h_{ij}) \leq \text{Val}(h'_{i'j})$ for any h_{ij} in $K_i \cap B_j$ and $h'_{i'j}$ in $K'_{i'} \cap B_j$. From this it follows that $P_\alpha \preceq_e P'_{\alpha'}$.

account of classifying *states of affairs* as right or wrong, while utilitarianism was concerned with classifying *actions*. Horty's *dominance ought* clearly goes a long way towards closing another gap as well: the one between deontic logic and decision theory.

Because of the weakness of dominance reasoning, however, Horty's account seems of limited value as a theory of choice. This chapter suggests how, with modest extensions, Horty's framework can move beyond dominance into the three main branches of the theory of decision: decisions under ignorance, decisions under risk and game theory. This leads to a motivational question: what is the point of trying to bring deontic logic "up to speed" if we already have a successful decision theory? I close by suggesting two main ways in which deontic logic provides return benefits for decision theory.

The first, noted at the outset of this chapter, is by offering rigorous analysis of foundational notions: causation, choice, counterfactuals and background states. That such analyses matter should be clear to anyone who has followed the history of decision theory as formulated by Savage, modified by Jeffrey, and re-formulated by causal decision theorists. For example, we claimed here that the *states* of decision theory are causal background contexts and provided an analysis of causal independence and background contexts within *stit* models. By contrast, Joyce (1999, 61) writes that *states* include all "aspects of the world that lie outside the decision maker's control". He tells us that future choices and events, if relevant to our present decision problem, must be incorporated into the background states for that decision. Now it is harmless to incorporate future choices and events into the background states if they are causally independent of the agent's present choice, but not so harmless if their future occurrence is contingent upon present choices. *Stit* frames take care of this automatically: histories belonging to distinct states at m must be *divided* at m . This rules out treating future choices or processes as constituents of states at m . Future chance processes must be incorporated into decisions via conditional chances for outcomes (as described in Sect. 6). To handle sequential choices in the *stit* framework requires something like Horty's *strategic ought* (2001, Chap. 7), which takes us beyond the present discussion.

As a second benefit, deontic logic offers a model for thinking about problems where decision theory and game theory cannot offer clear, uncontroversial solutions. One source of such problems is *infinite decision theory*, comprising decision problems in which an agent has to deal with infinite utilities, an infinity of possible acts, or both.³⁰ Some of these problems are genuinely paradoxical and have no clear solution. In other cases, however, there are clear prescriptions, yet decision theory is silent because there is no optimal act. Because deontic logic is concerned with the truth of obligation sentences $O[\alpha \text{ cstit}: A]$ even where A does not describe an act, it has the resources to offer advice in such cases.

One example of this kind, noted earlier, is the *greatest integer game*. Decision theory cannot recommend the choice of any particular integer, but our deontic logics tell us that $\odot[\alpha \text{ cstit}: A_n]$ and $O_m[\alpha \text{ cstit}: A_n]$ where A_n is the proposition " α

³⁰ See (Sorensen 1994) for examples.

chooses an integer larger than n ". As a similar example, imagine that α is a perfectionist attempting to finish a journal article. Suppose that α represents his position to himself as an infinite choice situation where A_0 stands for not submitting the chapter at all, A_1 for submitting the current version as is, and A_2, A_3, \dots for producing and submitting polished versions, each A_{n+1} slightly better than A_n . Suppose that all of the relevant utilities are bounded above by some fixed limit. In such a case, no act A_n is optimal. But our deontic logics still give us $\odot[\alpha \text{ stit}: \sim A_0]$ and $O_m[\alpha \text{ cstit}: \sim A_0]$, representing the obligation to submit the chapter.

Decision theory need not always be concerned about the metaphysical details of choice, or the precise characterization of the acts and background states needed to specify a decision problem. But at the boundaries of decision theory, where those details matter, *stit*-based deontic logic, made possible thanks to Belnap's rigorous analysis of agency, provides a wonderful resource.

Acknowledgments I acknowledge support from the Australian National University, in the form of a Visiting Research Fellowship.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution Noncommercial License, which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Belnap, N. 1992. Branching space-time. *Synthese* 92(3): 385–434.
- Belnap, N., M. Perloff, and M. Xu. 2001. *Facing the future: agents and choices in our indeterminist world*. Oxford: Oxford University Press.
- Broersen, J. 2011. Modeling attempt and action failure in probabilistic stit logic. Proceedings of the twenty-second international joint conference on, artificial intelligence, 792–8.
- Chellas, B. 1992. Time and modality in the logic of agency. *Studia Logica* 51: 485–517.
- Gibbard, A., and W.L. Harper. 1978. Counterfactuals and two kinds of expected utility. In *Foundations and applications of decision theory, of western ontario series in philosophy of science*, eds. C.A. Hooker, J.J. leach and E.F. McClennen, vol. 13, 125–162. Dordrecht: D. Reidel Publishing.
- Gillies, D. 2000. *Philosophical theories of probability*. New York: Routledge.
- Joyce, J. 1999. *The foundations of causal decision theory*. Cambridge: Cambridge University Press.
- Horty, J.F. 2001. *Agency and deontic logic*. Oxford: Oxford University Press.
- Kooi, B., and A. Tamminga. 2008. Moral conflicts between groups of agents. *Journal of Philosophical Logic* 37: 1–21.
- Levi, I. 1997. *The covenant of reason: rationality and the commitments of thought*. Cambridge: Cambridge University Press.
- Resnik, M. 1987. *Choices*. Minneapolis: University of Minnesota Press.
- Skyrms, B. 1980. *Causal necessity: a pragmatic investigation of the necessity of laws*. New Haven: Yale University Press.
- Skyrms, B. 1994. Adams conditionals. In *Probability and conditionals: belief revision and rational decision*, ed. E. Eells, and B. Skyrms. Cambridge: Cambridge University Press.
- Sorensen, R. 1994. Infinite decision theory. In *Gambling on God: essays on Pascal's wager*, ed. J. Jordan. Savage: Rowman and Littlefield.
- Spohn, W. 1977. Where Luce and Krantz do really generalize Savage's decision model. *Erkenntnis* 11: 113–134.

Nuel Belnap on Indeterminism and Free Action

Müller, Th. (Ed.)

2014, VI, 409 p. 42 illus., Hardcover

ISBN: 978-3-319-01753-2