What Must, Ought, Is Worth and May Not Be Done

Hamish Taylor
Department of Computing and Electrical Engineering
Heriot-Watt University

Abstract

The Praxis model of reasons for action takes choice among alternative courses of action to be a matter of what must, ought, is worth and may not be done. It envisages practical decision making systems being given their reasons for action as requirements, duties, rules of valuation, proscriptions and contingent circumstances in a knowledge base. The model's principles of selection and automated deduction then elicit from them what is to be done. Starting from local circumstances and rules specifying what must or ought to be done in various eventualities, initial options for consideration are deduced. Means end reasoning refines them into practicable alternatives. Alternatives to actions which must be done or actions entailing proscribed activity are excluded. Remaining alternatives are classified along incommensurable dimensions of value specified by relevance principles to bear on the decision context. Choice among alternatives is made using selection rules like the dominance principle. It states that alternatives that are more valuable in one way and no less valuable in any other are better ones. If one is better than any other, it is selected. If not then recourse is made to default selection rules. They provide reasonable ways for resolving conflicting rankings of alternatives along incommensurable value dimensions, for discriminating among indifferent alternatives, or for weakening conflicting requirements so that forbidden options become choosable.

Keywords: reason, practice, value, duty, proscription, axiology, deontology

1. Introduction

A rich variety of concepts and principles are open for use in reasoning about what to do. They provide more than one possible basis for developing a useful and valid model of practical decision making. Whereas Bayesian decison theory [6,8] begins with the relative desirability of consequences of actions open to an individual and his partial belief in such consequences happening upon performing these actions, this paper starts with the requirements, duties and proscriptions governing acting and the values yielded in performing alternatives, and develops a different kind of model of practical reasoning.

The Praxis model is intended for use in knowledge based approaches to automating practical reasoning. Reasons for decison making about what to do will be specified in the knowledge base in a propositional form, and depending upon their category serve as requirements (musts), duties (oughts), proscriptions (may nots), value relevance principles, canons of valuation assessment for kinds of action, or default rules for resolving what is to be done when no answer is obtained using the basic option selection process. Updating the knowledge base with contingent circumstances and deducing consequences from them and the stored reasons for

acting, according to the principles of the Praxis model, will answer the question of what is to be done.

The paper begins by outlining the Praxis model. It then contrasts the model with Bayesian decision theory. After that it develops the notions of requirement, duty, proscription and value in more detail to show how they can serve in an alternative model of practical decision making. It ends with some brief remarks on resolving blocked choices.

2. The Praxis Model

The Praxis model takes the basic question of practical reasoning to be

What is to be done?

Tactical, strategic and methodological questions about how it is to be done involving

agency	which agent(s) is to do it and how is answering the basic question to cause that
control	how is doing it to be steered and managed
methods	how are techniques and procedures to be employed in doing it
instrumentality	how are tools, aids and settings to be employed in doing it
scheduling	when and in what order are means and opportunity to be exploited in doing it

are beyond the scope of this paper, as are interaction questions concerning how to act in contexts where hostile or cooperative agencies are present and may intervene. Such questions are particularly complex where communication in the form of threatening, bargaining or reasoning with such agencies is possible. Also beyond this paper's scope are questions of how exigencies of practical decision making like limited time, resources, and knowledge to deliberate with, and limited agency to enact decisions with, affect deciding what is to be done. This paper leaves these questions to one side. It addresses the issue of what is to be done in terms of reasons for and against acting as four separate questions:

Question	*Type of Reason*	*Study*
What must be done?	requirement	Deontology
What ought to be done?	duty	Praxeology
What may not be done?	proscription	Deontology
What is worth being done?	value	Axiology

The decision circumstances together with general requirements specifying what must be done in various eventualities may imply that some actions must be done now. More probably the same circumstances and general praxeological principles expressing duties, needs and interests of the decision maker only imply what ought to be done now. Both requirements and duties can arise out of a role (e.g. professional expert or assistant) or be the raison d'etre of an automaton (a Mars exploration robot) or be duties attached to the purpose of a decision support system (advising about financial affairs). Needs and interests arise out of what it means for the decision maker to have itself as an agent to enact some of its decisions. Needs and interests specify what ought to be done to maintain that agency's well-being, power, and potential to do things.

Whereas requirements exclude their alternatives, these duties, needs and interests are likely to generate conflicting options for the decision maker to consider in deciding what is to be done. These options are then confronted with contingent circumstances in means end reasoning, in ways not discussed in this paper, to elicit practicable courses of action, and then regimented to form a set of alternative things to be done. The identification of alternatives enables courses of action which must be done to exclude their alternatives.

The third question of what may not be done is now addressed to ensure that alternatives involving forbidden actions are also excluded from consideration. The role of these considerations is to set limits on certain possibilities of practical decision making. While some practical decision making might not be subject to any requirements or proscriptions, critical matters such as security and safety invite decision making in terms of what must be done and what may not be done in those circumstances.

The Praxis model offers two variations to cope with conflicts among requirements and/or proscriptions. The hard Praxis model imposes an integrity constraint on acceptable sets of requirements and proscriptions which precludes the deducibility for any circumstance of conflicting musts or may nots. This may be achieved by global consistency checking or by prioritising the bearing of conflicting requirements. The soft Praxis model imposes no such integrity constraint but handles conflicts between rival mandatory alternatives or between a mandatory course of action and a proscription of an action entailed by it by equitably downgrading the force of relevant requirements to lesser force ought to be dones, and the force of proscriptions to attaching an apt strong negative value to performing the action.

In the Praxis model if only one thing must be done, then that thing answers the question of what is to be done. If more than one thing may be done, then the fourth question of what is worth being done is addressed. Principles of relevant values specifying which kinds of values are pertinent to the decision context are applied and the remaining alternatives are evaluated using canons of valuation along each such dimension of value. Alternatives are then compared using evaluative principles. They specify sufficient conditions over the values of alternatives A and B under which

A is a better alternative than B

These evaluative principles order alternatives partially or totally. If they order them totally - i.e. there is one alternative which is better than any other, it becomes the alternative which is to be done. If there is no single course of action which is better than every other alternative, then the Praxis model allows default selection rules to be applied in turn over the elicited alternatives to try to produce a totally ordered ranking of alternatives. Various such principles are proposed in a later section. None of them can be defended as a basis for making a best choice, but each has some merits as a way of making a reasonable choice among alternatives. They are applied independently or iteratively until the alternatives are totally ordered in merit, or value indifference is sufficiently established among a set of alternatives, all of whose members are better than all non-members. In the first case a single answer is supplied to the question of what is to be done, and in the latter case any member of the indifference set is an answer to the question.

3. Bayesian Decision Theory

The Bayesian approach to normative decision theory [6,7,8,12] analyses practical decision-making as follows. A decision is a choice among a set of pairwise exclusive and exhaustive alternative things to do.

$$A_1, \ldots, A_n$$

Each alternative A_j has an exhaustive and exclusive set of consequences

$$C_1, \ldots, C_m$$

Each consequence C_i has a probability Pi between 0 and 1 of coming about if A_j is done, and that consequence has a finitely valued utility U_i to an agent *I* of happening. The merit of choosing an alternative A_j to the agent *I* is the sum of the expected utility E_j of all its consequences:

$$E_j = \sum_{i=1}^{m} P_i \times U_i$$

The Bayesian choice rule is to choose a course of action Ak which has a maximal expected utility

$$E_k = \max\{E_1, \ldots, E_n\}$$

The possibility of ties means this choice may not be unique.

The Bayesian model embodies four striking characteristics

- consequentialism
- relativism to alternatives
- value monism
- synthetic subjectivism

Alternatives are only evaluated in terms of their consequences. They are not valued in terms of the merit of the action itself. Furthermore choice is always relative to alternatives. Whether an alternative gets chosen or rejected depends only on whether its expected utility is more than any other or less than another alternative. It never depends only on the reasons relating to the alternative itself. A third characteristic is that only one value utility is used to assess the merits of an alternative's consequences. Multiple dimensions of worth for assessing whether to choose an alternative are not recognised. Furthermore this value is not one belonging to states of affairs or things but is synthesised from individual preferences using betting measures [7] or other means [6]. Thus the approach is subjectivist and delivers conclusions which are only valid for an individual with those subjective probability and utility assignments. By contrast the Praxis model embraces

- non-consequentialism
- singularism about alternatives
- value pluralism
- naturalistic intersubjectivism

The Praxis model allows alternatives to have reasons for doing them which are not stated in

terms of their consequences. Doing that alternative may be a requirement or duty of a role, or be valued like courage or integrity for the sake of its own performance. The Praxis model also allows alternatives to be chosen or rejected solely by their attached reasons irrespective of alternatives. An alternative which must be done or one which may not be done can be chosen or rejected on its own merits irrespective of what else should or must be done. The Praxis model is also based on accepting that there are many different values, which cannot be reduced or measured by a single one. This ineradicable value pluralism creates the problem of resolving conflicting valuations among incommensurable values for reasoned choice. The Praxis model encompasses methods for its reasonable management and views doing so as central to most interesting practical decision-making. Lastly the Praxis model endorses a naturalistic intersubjectivism. It treats value ascriptions as claims about the nature of things and states of affairs. It is neutral between supposing possession of a value is an objective property of a thing and supposing that it is intersubjectively constituted within a normative framework whose scope of validity is conditioned by socio-cultural, historical and anthropological facts. This stance empowers the Praxis model to deliver conclusions with intersubjective validity about what is to be done.

4. Reasons for Acting

Natural language expresses possibilities using a family of modal terms. They include “must”, “ought”, “may”, “might”, “can”, “could”, “should”, “has to”, “is necessary” and “is possible”. These terms express different kinds and strengths of possibility [4] appendix B. Three strengths of these modalities can be distinguished.

strong terms:	*must, has to* and *is necessary*
intermediate terms:	*ought* and *should*
weak terms	*may, might, can* and *could*

Thus the propositions

Kate must eat her pudding

Kate ought to eat her pudding

Kate may eat her pudding

express a progressively weaker set of propositions regarding what Kate has reason to do. The different strengths provide the Praxis model with different categories of reasons which constrain practical decision making in quite different ways.

5. Deontology

The strong form “must” and the negated weak form “may not” provide an intuitive idiom for expressing necessities for acting. It is natural to seek to formulate binding instructions which practical decision makers must obey and to specify courses of action which they may not perform. Thus codes of conduct are formulated for organisations like the British civil service, which specify what civil servants must or may not do. They cover matters like forbidding receiving gifts from persons with whom civil servants conduct their official business. In a similar way it is natural to seek to specify binding instructions for machines. The science fiction writer Asimov makes a good case for doing this with his three Laws of Robotics [2]

First Law	A robot *may not* injure a human being, or, through inaction, *allow* a human being to come to harm.
Second Law	A robot *must* obey orders given it by human beings except where such orders would conflict with the First Law.
Third Law	A robot *must* protect its own existence as long as such protection does not conflict with the First or Second Laws.

While it is possible to cavil at Asimov's formulations of his Laws, Asimov is persuasive in arguing that robots (and decision support systems) need to be subject to such binding restrictions which admit of no exceptions. If the decision making of machines is to be trusted, absolute guarantees are needed that certain matters such as allowing injury to human beings or letting itself be destroyed, won't get permitted in some complex trade off of reasons for and against certain courses of action. Although it is somewhat ironic that many of Asimov's pieces of robot fiction turn on just those exceptional cases where violations of these Laws would make a certain sense, this argument still seems persuasive. By giving such considerations overriding priority in a sound deductive framework, it should be possible to frame safety constraints for practical decision making by machines and prove that such machines will never knowingly violate them. This would be useful in automated reasoning applications like control of vital monitoring equipment or dangerous engineering plant where maintenance of safety is a critical issue.

Right away Bayesian Decision Theory becomes an inept vehicle for articulating the application of such practical precepts to guide action. At heart it is a mechanism for trading off any course of action against any alternative depending on their relative costs and benefits. No course of action is deemed forbidden or mandatory irrespective of alternatives. So long as all utilities of consequences are finitely valued, it remains possible that the expected utility of any action may be higher than any of its alternatives.

A more apt vehicle for expressing considerations like Asimov's three laws is Deontology. Deontology is concerned with what must, may or may not be done. Deontology is often misleadingly described as the study of duty. This conflation of duty with what one must do has been obscured by the idea that duty is somehow compelling. However, it is not plausible to suppose that duties are always compelling, in the sense of overriding other considerations for acting. Intermediate strength terms in the deontic modality like "ought" are more apt for expressing the claims of duty than terms like "must".

5.1 Deontic Logic

Deontic logic [13] attempts to capture the properties of what must or may not be done. It uses a modal operator **P** which qualifies propositions.

P p *means* it is permissible that p

Using Davidson's method for rendering action sentences [3] which gives them an extensional representation by introducing an extra argument for the action's event, the first part of Asimov's First Law of robotics can be rendered as follows

$\neg$**P** (all event e) (all thing x) (all person y) (robot(x) $\supset$ injure(x, y, e))

The relation between what is mandatory and what is permissible is intuitively captured by the idea that something must be done if and only if it is not permissible that it may not be done.

Thus Asimov's Third Law of robotics without its exclusion clause is

$$\neg\mathbf{P}\,\neg(\text{all thing x})\ (\text{all event e})\ (\ \text{robot(x)} \supset \text{protect(x, x, e)}\ )$$

Standard formulations of deontic logic regard

$$\neg\mathbf{P}\,\neg\ \text{p}\ \&\ \neg\mathbf{P}\ \text{p}$$

as a contradiction. The significance of this for the pragmatics of using this logic to articulate reasoning with permissions, is that the set of what is mandatory and what is forbidden must be formulated to be free of contradictions. This approach is adopted by the hard variation of the Praxis model. However, it is a tough requirement to meet without introducing ad hoc measures.

5.2 Conflicting Demands

Suppose that a Mars explorer robot has the following mandatory requirement in its knowledge base.

A report must be made to the command base on Earth every day.

and the following proscription in its knowledge base.

The backup generator may not be used to power the main radio transmitter.

The latter we can suppose has been put in to avoid the chance of damage being caused to the backup generator by the heavy power drain of using the main transmitter. Now it may happen that means end reasoning determines that, given the current state of the primary generator, the only way to power the main radio transmitter to make the daily report to the command base on Earth is by using the backup generator. Furthermore means end reasoning may determine that the daily report can only be made to the command base on Earth by using the main transmitter. The Mars explorer robot is now in the impossible situation of being subject to contradictory demands to use and not to use the backup generator.

Asimov addresses this kind of problem by prioritising the application of his Laws of Robotics. Undoubtedly this method could be made to work for the case above by attaching priorities to requirements and proscriptions. However, the worry is that such a method is blindly powerful. It is too easy to prioritise all musts and may nots, but it is much harder to be sure that such prioritisation will handle conflicts in a sensible way.

A different approach tries to dissolve the incompatibility by equitably downgrading inconsistent demands that generate the conflict for this local part of the decision making process. This is the approach of the soft variation of the Praxis model. The requirement that a report must be made to the command base on earth can be downgraded to a duty that a report ought to be made. At the same time the proscription that the backup generator may not be used to power the main transmitter can be waived and the force of its injunction transmuted into a relevant valuation which would give strong relevant disvalue to performing that action.

It can be argued that allowing requirements and proscriptions to be downgraded in this way undermines the point of having them at all. If a requirement or a proscription can be waived then it fails to fulfill its function of necessitating performance or non-performance of a course of action. In defence of the soft Praxis model it can be replied that waiving a requirement or proscription does not preclude the nature of the consideration from having considerable force in determining what is to be done. Furthermore such considerations are only waived when

there are contradictory necessities governing what is to be done. Necessities are not discarded lightly on the soft Praxis model.

Neither the hard nor the soft Praxis model adopts an ideal solution. Each has its merits and demerits. It would be invidious to prefer one solution to the other. So both are countenanced as acceptable variations of one model.

6. Praxeology

Whereas the strong modality of requirements and proscriptions expresses some useful necessities of acting, the degree of its own modality renders it unable to articulate genuinely conflicting reasons for acting. Intermediate modal terms such as "ought" offer a better idiom for articulating and dealing with conflicting reasons for acting. Because "ought" is a term of intermediate modal strength, there is no contradiction in saying of someone that he ought to do *F* and ought not to do *F*. The apparent conflict can be explained away by noting that each use of "ought" will be relative to a different reason. By contrast there would be a contradiction in using the strong modal term "must" and saying of someone that he must do *F* and must not do *F*. The point of using the strong modal term "must" is to exclude the possibility of his not doing *F*. No relativisation to reasons saves the statement from contradiction.

Two different forms of relativisation to reasons by "ought" can be distinguished:

prima facie	relative to a reason;
all things	considered relative to all relevant reasons.

Although a person may have several alternative things he ought prima facie to do, he can have only one thing he ought to do all things considered. Ought propositions can also be qualified with regard to a set of considerations. Thus the forms ought morally, ought legally and ought as a matter of etiquette can be distinguished. Sometimes the qualification morally to a use of ought is intended but is not made explicit. This can be detected by asking whether in the decision context what ought to be done is what ought to be done morally or what ought to be done all things considered.

Another relevant distinction is that between what ought to happen or be the case and what ought to be done. Although something ought to happen or be the case, it does not follow that it ought to be done. What ought to happen may not be open to being achieved by action, or it may only be open to being achieved by collective forms of action which are not open to any single agency (collective or otherwise) to effect. This point is really underpinned by the principle

P1 What ought to be done must be within the power of some relevant agency.

This idea is often expressed by the slogan "ought implies can". If what ought to happen is not within any agency's power to effect, then it is not the case that it ought to be done.

It seems plausible to suppose that practical reasoning conclusions about what ought to be done, can only be validly derived from premisses containing a praxeological principle about what ought to be done. Such conclusions cannot be validly derived from premisses containing no ought premisses at all. This principle is often summed up in the slogan "*an ought cannot be derived from an is*". Hume is usually credited with authorship of this useful principle in his *Treatise of Human Nature* Book III, part I, section I [5] p.469.

This principle is useful as a heuristic in searching for suppressed premisses in practical reasoning. In a rule such as

IF patient P has an inflammatory joint disease &
oral drugs are suitable for patient P &
patient P is not suspected of having a peptic ulcer
THEN patient P ought to take aspirin orally

it suggests that a reason why the patient ought to do something has been suppressed, and should be brought out. After all the patient's taking aspirin is only a means to something else, his regaining his health. So that should be brought out explicitly. In this way the principle motivates the development of a model of practical reasoning where all of an agent's ultimate "ought" reasons for acting are made manifest as a set of praxeological principles for acting. In a natural sense some of these might be called the duties of an agent.

With human beings these can arise from an agent's roles, such as father, husband, friend, official, citizen etc., or they can be duties of codes of conduct which govern the agent's behaviour, such as ethics, local custom, honour code, religious convention etc.. With an automated practical reasoning system duties will arise from its raison d'etre, what its job is. Needs and interests of an agent will also serve as a source of possible actions which an agent ought to perform.

A corollary of the is-ought principle is that means end reasoning propagates oughts from ends to means. Clearly a doctor ought to treat patients of his who consult him. Furthermore if such a patient upon consulting his doctor is found to need to be prescribed a drug, then that doctor ought prima facie to prescribe his patient that drug. Here the "ought" attaching to the doctor's duty has by means end reasoning been propagated to the means for accomplishing the end.

Agents can have both positive and negative duties. Two such duties of a doctor might be:

> A doctor ought to inform the authorities upon discovery that a patient of his has a notifiable disease.
>
> A doctor ought not to disclose details of a patient's medical history to unauthorised persons.

In the Praxis model of practical reasoning positive duties provide broad kinds of things for an agent to consider whether they are to be done. However, negative duties cannot do the same thing. They count against performing certain kinds of action. In order to be brought into play, an action of the type they count against has to be licensed by a positive duty. Hence negative duties are really either constraints on duty or part of the evaluation of duties and not part of the generation of alternatives for consideration. Thus a negative duty really has to be represented either as something forbidden in the proscriptions part of the Praxis model or as having negative value in the valuation part of the model, and not in the duty part we are considering here. A complication, which we shall not discuss, is that it is not always easy to distinguish positive from negative duties. A positive duty to keep a promise might also be represented to be a negative duty not to break a promise.

The properties of the relevant form of ought can now be summarised as follows. Its logical form is

> ought(Agent, Deed, Reason)

An ought of reasons is relative to an agent, a deed, and a reason. Agents are usually individuals but can also sometimes be collective entities like organisations, teams etc.. Deeds are actions of the agent. Reason is a single or a set of reasons for performing the deed. Some relate the Deed as means to a further end in acting. Others give a description of the type of end that the Deed is. Such oughts are usually applied by conditions, which determine the circumstances, context and time, under which the reason applies to the agent. The following conjunction relating an agent A and a deed D is not a contradiction.

$$(\exists\ R1)(\exists\ R2)\ \text{ought}(A, D, R1)\ \&\ \neg\text{ought}(A, D, R2)$$

Two properties observed by "ought" are:

$$\text{ought}(\text{agent}, \text{deed}, \text{reason}) \supset \text{can\_do}(\text{agent}, \text{deed})$$

$$\text{ought}(He,Deed,Reason)\ \&\ \text{means}(He,Act,Deed) \supset (\exists\ R)\ \text{ought}(He,Act,R)\ \&\ R = \text{in\_order\_to}(Deed,Reason)$$

The first implies the Deed is within the agent's general power to perform. The second states that that if an Act is a means for an agent to perform a Deed and He should do that Deed for a Reason, then He should perform that Act in order to do that Deed for that Reason.

7. Value Theory

Practical decision making is not only involved with doing what must be done, avoiding what may not be done and doing as much of what ought to be done as can be done. It also involves selecting among alternative things which ought to be done. This involves choosing better alternatives over worse alternatives on the basis of their relative merits. In an impersonal model of practical reasons, this means in terms of their relative worths. Thus the Praxis model requires the means to assess the value of alternatives and judge their respective worths. To get this we must turn to value theory [9 chap 5, 10].

A value can be characterised as a beneficial quality of an action or a state of affairs, which would give a reason inter alia for performing it or for bringing such a state of affairs about. Values are values because their attainment and maintenance contributes to well-being and the good life. However, they are not subjectively constituted by individual attitudes to such matters. They have an interpersonal standing, and in this sense are impersonal [10] p. 11.

Trying to incorporate into the Praxis model a mechanism to adjust the worth of a value by whether and how much it is valued by a particular individual would be to abandon the impersonal standpoint. It would render all adjusted valuations idiosyncratic. The Praxis model is better off remaining as impartial as it may, so that its valuations remain grounded in a contemporary framework of evaluation which has as much intersubjectively constituted authority as it can get. Careful choice of apt value concepts should ensure that the right standards of valuation get applied to appropriate decision contexts.

Values can be either instrumental or intrinsic. Instrumental values are valuable because they contribute to the realisation of other values, whereas intrinsic values are valuable in their own right. Intelligence would be an instrumental value in making a good employee whereas truth would be an intrinsic value in theoretical inquiry. Values can be classified in various ways [10]

Classification Dimension	*Examples*
subscribers	personal, professional, or national values
domain of application	bravery in humans, purity in things, justice in societies
nature of benefit	economic, aesthetic, or moral values
purposes at issue	exchange, food, or persuasive value
subscriber-beneficiary relation	egocentric, disinterested values

Values can have either an ordinal value scale

unhappy equable happy

or a cardinal value scale like cost (£1000s). These value scales can be bounded or unbounded. Most values are bounded in having limiting degrees on their scale at both ends. Thus there is no greater degree of integrity than full integrity and no lesser degree of lack of integrity than complete lack of integrity. However, some values like cost are unbounded in having no limit to their scale at their disvaluable end. Some people have believed that the dimension of good and evil is unbounded at both ends. Namely, that however evil or good a state of affairs is, more evil or more good states of affairs might exist than it.

Values can be monopolar in the sense of having only a single dimension of worth, or be bipolar and have both a positive dimension of value and a countervailing negative dimension of disvalue. Bipolar values have a median or neutral point on their scale which is the possession of no degree of disvalue nor of value. Grace is best construed as monopolar in having the following scale

graceless partly graceful graceful

whereas beauty is a bipolar value

ugly plain beautiful

Ordinal scales of value can usually be partitioned in coarser and finer fashions. Thus the scale of beauty above might be refined into a finer scale of gradations as follows

very ugly ugly plain beautiful very beautiful

Important concepts in generating bipolar dimensions of value are good and bad. The attributive form “a good X” expresses the idea that the X is valuable for use as or in the role of an X. The related form “a bad X” expresses the polar opposite. These concepts are useful for creating concepts of (dis)value which sum up various components aspects of value. The scale

bad knife indifferent knife good knife

sums up various instrumental values like robustness, sharpness, easiness to wield etc..

The dimensional nature of value is underpinned by the following principle

P2 Something is a value only if other things being equal it is more valuable to exhibit or realise more of it and less valuable to exhibit or realise less of it.

This dimensionality of value principle is important for screening putative values to see whether they are really values. The principle shows that a parameter like a house’s size is not an instrumental value in making a good house. While greater size in a house often makes it more valuable other things being equal, there is also such a thing as a house being too big to be

a good house. A good house should be of an apt size. Aptness of size not size is the pertinent instrumental value. Other things being equal a house of a more apt size is always a better house and a house of a less apt size is always a worse house.

Flexibility in value gradations shows its usefulness, because choice can and should be more or less sensitive to matters of degree along dimensions of value. An employer choosing a front office receptionist should recognise that only whether a prospective employee is of good appearance or not is germane to the choice and that any more discriminating interest in exactly how attractive he or she is is inept. So long as the prospective employee is at least of good appearance that is all that matters with respect to that scale of worth. However when considering diligence, the employer should be more discriminating and be interested in making reasonably fine distinctions between degrees of diligence which would distinguish between being moderately, quite, rather, very and extremely diligent.

Values are related in various ways to action. Actions can manifest, create, destroy and maintain the existence of values [9] chapter 5. An important distinction here is that between performance values, state values and production values.

performance value	gracefulness, truthfulness, courage, reasonableness
state value	happiness, welfare, justice, truth

Performance values are transiently manifested during action. State values belong to objects, events or states of affairs. Some values like beauty can be both performance and state values. They can be transiently manifested during performances such as dancing and belong to objects in the world such as a statue. A third kind of value is the ability of an action to produce a state value. This kind of value of the action will be called its production value. The Praxis model is concerned with both the performance and the production values of alternative courses of action. By contrast Bayesian decision theory focuses on production values and ignores performance values.

A key feature of value is that there is an essential plurality of values. There is no plausible way of accounting for the worth of each value using some single category of value. There may be ways in which values can be related to each other. Thus instrumental values can be related to intrinsic values and intrinsic values may complement each other in organic unities [9] chapter 5. Also some partial rankings of values look plausible e.g. harmfulness to humans rates as more important than gracefulness. But even these value rankings don't determine how much gracefulness is outweighed by how much harmfulness. The existence of a plurality of values has far reaching significance. It means that the question of which of two alternatives is the better one often has no right answer. A course of action A may exhibit or realise more of the value V and less of the value W. An alternative B may exhibit or realise more of the value W and less of the value V. Thus the following relationships may hold

$$V(a) >_v V(b)$$

$$W(b) >_w W(a)$$

asserting that A's V value is greater than B's V value, but that B's W value is greater than A's W value. However, the different subscripts of > show that different ordering relations are in question. If V and W are distinct values, there is usually no way in which the differences between these values (if they are definable)

$$U(A\text{-}B) = V(A) - V(B)$$

$$U'(A\text{-}B) = W(B) - W(A)$$

can be rendered in some common value U and related by an ordering relation "$>_u$" in a fashion such as

$$U(A\text{-}B) >_u U'(A\text{-}B)$$

Hence, there is no way that the overall relative value of A to B can be determined. Frames of consideration of practical decision making which exhibit this kind of structure will be said to exhibit a conflict of values. They are encountered very often. However, they do not render reasoned choice intractable as will be seen.

However, not all practical decision making encounters conflicts of values. In many decision contexts the existence of a plurality of values can be admitted and yet a better choice be identified. The key to doing this is the dominance principle that better alternatives exhibit or realise more value then worse ones.

P3 If a course of action A is worth more in at least one relevant respect of value and is not worth less in any other relevant respect of value than an alternative course of action B then A is a better alternative than B.

This principle does not say that either alternative should be chosen. It only identifies the better alternative. The existence of a plurality of values requires it to be framed in this form, because there is no general way for measuring the relative worth of one grade of value in terms of another (c.f. Richards' teleological dominance principle in [11]).

7.1 Example of Evaluation

All significant relative appraisal involves a significant degree of interpretation. In the Praxis model this will proceed along several different dimensions of value. Relevance principles will clarify which values are relevant and at what level of discrimination, and a process of appraisal will be needed to transform the manifest parameters in which the choices present themselves into the relevant categories and grades of value and disvalue. It is not to be expected that there will be a precise theory of value measurement to govern how this should be done.

When choosing to buy a house to live in, options can be considered in terms of how big the house is, how near various main roads and railway stations it is, what the house and its surrounds look like, what its views are, how many bedrooms it has, what its cellars, kitchen and living rooms are like, whether there is anything wrong with it structurally, how much it costs etc.. These factors can be assessed by inspection or from a surveyor's report. Evaluation now requires that these parameters be translated into grades on relevant scales of value.

Praxeological and means end reasoning might deliver three alternative options in a house buying choice

Ashgrove ought to be bought.

Bellevue ought to be bought.

Cornhall ought to be bought.

Features of these three options might be

House Name	Bedrooms	Transport Links	Situation	Modern Fittings	Survey Report	Price
Ashgrove	5	1km to station	countryside	all mod cons	OK	£88K
Bellevue	4	5km to M25	leafy suburbs	old but OK	excellent	£78K
Cornhall	5	3km to station	terrace house	all mod cons	minor damp	£82K

In order to render these into the Praxis model, it is necessary to determine what the relevant values are, and what their grading scales are. The following appraisal values are plausible scales for assessing the merits of these three choices

size aptness	how well the house size fits the need
accessibility	how well the house has good transport links
amenity	how well the house's situation conduces to living there
utility	how well the rooms and layout serve the requirements
soundness	how structurally sound the building is
cost	how much the purchase of the building and grounds is

Thus the model will need to contain a relevant value principle such as the following:

> Relevant values for choosing a house to purchase and live in are size aptness, accessibility, amenity, utility, soundness, and cost.

Perhaps this principle does not capture all relevant values but it suffices for illustration purposes. The model also needs to contain grading scale information such as the following:

Value	**Grading Scale**			
size aptness	poor	moderate	good	very good
accessibility	poor	moderate	good	
amenity	low	middling	high	very high
utility	poor	moderate	good	very good
soundness	unsound	sound		
cost	cardinal scale in units of £15 000			

This information can also be rendered in propositional form in the model's knowledge base as

> The grades on the scale of accessibility in choosing a house to purchase and live in are in ascending order - poor, moderate and good.

When appraisal is finished, the assessment of the values of the house choices might be as fol-

lows.

Name	Size Aptness	Access	Amenity	Utility	Soundness	Cost
Ashgrove	good	good	very high	very good	sound	£75-90K
Bellevue	moderate	good	high	good	sound	£75-90K
Cornhall	good	good	middling	very good	sound	£75-90K

The appraisal places each alternative on each relevant dimension of value using a relevant grade of discrimination. When the process is complete, the formal principle of choice can be applied. In our example the most expensive house “Ashgrove” comes out as the dominant choice. If there is no other alternative then all things considered Ashgrove ought to be bought.

7.2 Formal Theory of Value

This example can be formalised as follows:

Relevant Vocabulary

1) 3-place relational expressions buys(x, y, e), $x >_v y$ and $x =_v y$
2) variables for events e, f, g, for values v, w, for houses x, y, z, for persons i, j
3) 2-place *better alternative* connective $\Rightarrow$
4) house constants A, B, C
5) value constants Size, Access, Amenity, Utility, Sound, Cost
6) event constants X, Y, Z

The rules of well-formedness are the same as for the predicate calculus where event, house, person, and value variables are taken to be individual variables, the better alternative connective is taken to be a two place connective between relational expressions. The 3-place relations $x >_v y$ and $x =_v y$ stand for whichever 2-place value ordering relation the value variable v is interpreted as.

The interpretations of the terms are all straight forward. The constants A, B and C represent the three houses beginning with those letters. The constants X, Y and Z represent the three possible events of buying the houses A, B and C. The values constants are abbreviations of the relevant value names. The value ordering and equality relations are indexed by abbreviations of the relevant value dimensions.

The value ordering relations $>_v$ are irreflexive, asymmetrical, and transitive, and the value equality relations are transitive, reflexive and symmetrical. These assumptions are stated as A1-A6.

Assumptions of Theory

A1 $(\forall e) \neg e >_v e$

A2 $(\forall e)(e >_v f \supset \neg f >_v e)$

A3 $(\forall e)\,(\, e >_v f \,\&\, f >_v g \supset e >_v g\,)$

A4 $(\forall e)\,(\, e =_v f \,\&\, f =_v g \supset e =_v g\,)$

A5 $(\forall e)\; e =_v e$

A6 $(\forall e)(\, e =_v f \supset f =_v e\,)$

A7 expresses the relationship between value equality and value ordering.

A7 $(\forall e)(\, e =_v f \supset \neg\, e >_v f\,)$

The detailed valuations of the theory are stated in A8-A19.

A8 $X >_{Size} Y$

A9 $Z >_{Size} Y$

A10 $X =_{Access} Y$

A11 $Y =_{Access} Z$

A12 $X >_{Amenity} Y$

A13 $Y >_{Amenity} Z$

A14 $X >_{Utility} Y$

A15 $Z >_{Utility} Y$

A16 $X =_{Sound} Y$

A17 $Y =_{Sound} Z$

A18 $X =_{Cost} Y$

A19 $Y =_{Cost} Z$

These assumptions state the relationships along each dimension of value. These events X, Y and Z are the three (potential) acts of purchase. They are related to buying by the auxiliary assumptions

E1 $(\exists i)(\exists e)(\, buys(i, A, e) \,\&\, e = X\,)$

E2 $(\exists i)(\exists e)(\, buys(i, B, e) \,\&\, e = Y\,)$

E3 $(\exists i)(\exists e)(\, buys(i, C, e) \,\&\, e = Z\,)$

In order to bring overall merits into the picture a version of the dominance principle is needed. It is stated as two propositions A20 and A21.

A20 $(\forall e)(\forall f)((\exists v)(\, e >_v f\,) \,\&\, (\forall w)\,(\, \neg\, f >_w e\,) \supset e \Rightarrow f\,)$

A21 $(\forall w)(\, w = Size \vee w = Access \vee w = Amenity \vee w = Utility \vee w = Sound \vee w = Cost\,)$

These assumptions A1-A21 suffice to prove the conclusions C1 and C2.

C1 $X \Rightarrow Y$

C2 $X \Rightarrow Z$

which establishes that buying Ashgrove is a better choice than either alternative. Since these are all the alternatives there are, buying Ashgrove is the thing to be done all things considered.

7.3 Risk and Valuations

The acknowledgement of a fundamental plurality of values and the admission that most values have only an ordinal and not a cardinal scale, deprives the Praxis model of the ability claimed by Bayesians to measure the relative merits of any set of alternative courses of action. However, this kind of model can still handle risk. Risk affects the worth of production values possessed by a course of action. The basic impact of risk on value is embodied as follows:

P4 If two alternatives A and B have chances P_a and P_b of producing a consequence of grade G of value V and $P_a > P_b$ then A has greater production value of V in that respect than B does.

This risk principle can be used directly as a basis of choice where risk is involved. Given the two alternatives of accepting a 50% chance of getting a £1 or a 30% chance of getting a £1, the 50% chance of getting a £1 has a greater production value of the same monetary value. Hence if monetary production value for the chooser is the sole value relevant to the choice, the dominance principle implies that the 50% alternative is the better one. However, given a choice between the alternatives T1 and T2

T1 50% chance of getting £1

T2 30% chance of getting £2

the risk principle is of no help in valuing the alternatives. It would be nice if some valid general principle for discounting production value in terms of risks could be established. Such a principle would specify the degree for discounting a production value of an action in terms of the chances of its valued consequences being realised. A consequence of such a discounting principle for a particular value might be the rule R:

R The monetary production value of a course of action is equal to the numerical product of the probability of the course of action producing monetary value with the numerical value of the money.

This rule R presupposes numerical measures of probability and money. Under some obvious assumptions it would allow one to value the alternatives A1 and A2 as having £0.5 and £0.6 monetary production values. If monetary production value for the chooser was the only relevant value in the choice situation, then it would enable one to identify A2 as the better alternative. However, it is highly doubtful that there can be any valid general discount principle which would entail a rule such as R. There is nothing intrinsic to the idea of valuation and risk that would justify such a principle. Rules such as R are highly dubious anyway because there is little reason to agree that there is no difference in monetary production value between the alternative choices D1 and D2.

D1 100% chance of getting £10

D2 50% chance of getting £20

The Praxis model operates with no general principle for discounting production values in terms of risk. It does this because under its assumptions it is not plausible to suppose that such a principle exists. However, rules such as R can still figure in the model as auxiliary postulates of evaluation or canons of assessment in order to determine how to grade the production value of a course of action. It is up to the user of the Praxis model to ensure that such use of auxiliary postulates will result in the right grades of production value being accorded to alternatives with

only a chance of realising valued consequences.

8. Blocked Decision Making

The Praxis model makes a sharp separation between straight forward applications of the model of practical reasoning in accordance with praxeological, deontological and axiological considerations, and use of the default choice procedures. Straight-forward applications of the Praxis model are intended to yield uncontroversially valid results. So long as either one thing must be done or there are a set of permissible alternatives such that

a) they ought to be done

b) they are correctly valued by relevant values at the right degree of discrimination

c) one is a better alternative than any other

then it is right that that thing be done. Where these assumptions do not all hold, the model cannot resolve the issue of choice uncontroversially. Nonetheless good scope remains for resolving the issue of choice in a reasonable fashion.

Where there is more than one permissible alternative that ought to be done and no alternative is better than every other, then no alternative is the best. However, some alternatives may still be better than others. Without loss of generality the set of alternatives can be reduced to those alternatives that no alternative is better than. This set of alternatives will be called the undominated set. By hypothesis there will be two or more members of this set. Now members of this set will either all have the same grades of the relevant values or not.

same valued alternatives	*value indifference*
differently valued alternatives	*value conflict*

The set of alternatives will either be in a state of value indifference or in a state of value conflict. A state of value conflict does not prevent a strict subset of the alternatives from being in a state of value indifference. The Praxis model requires that choice be made in both these circumstances, because each alternative expresses something that ought to be done. Not choosing would mean that less of what ought to be done would be done than otherwise. A state of value indifference can be handled either by random choice or by reassessing the way the alternatives are valued.

Random choice is apposite for the logical donkey equidistant between two identical piles of hay, because the alternatives by definition are exactly alike. This situation is not typical. Usually some quite plausible changes in ways of valuing alternatives will alter valuations among the set of alternatives. Each of them gives an opportunity for a dominant alternative to emerge. For this reason random choice looks unappealing as a way of resolving a state of value indifference. Using such a method evades looking for some reason for showing that one alternative is more worthy than another.

Reassessing the way alternatives are valued also provides a good way of resolving value conflicts. There are many possible methods for reassessing the set of valuations. Two of them are:

More Discrimination	making the valuation of alternatives more precise
Less Discrimination	making the valuation of alternatives less precise

In terms of the house buying example, it would be possible to increase or decrease the number

of grades of relevant values. Thus the utility scale could be given five or three gradations rather than four. The cardinal scale of cost could also have its granularity increased or diminished. Thus the units of significance for cost could be changed to £10,000 steps or to £20,000 steps. Clearly the less grades of values used, the more likely it is that a dominant alternative or a state of value indifference will emerge. Conversely the more discriminating valuation is, the more likely that the alternatives will be put into a state of value conflict. Nonetheless it is still possible to eliminate value conflicts by introducing more grades of values.

Let it be supposed that two alternatives A and B are graded in one order by the value V and in the other order by the value W. Furthermore let it be supposed that they are graded equal by all other relevant values. Lastly let it be supposed that A and B were graded as *slightly* and *moderately* valuable in terms of V. Introducing a new intermediate grade of V between *slightly* and *moderately* called "*mildly*" might result in both A and B being reclassified as *mildly.* This would dissipate the value conflict, and result in one alternative emerging as the dominant one.

Altering the scope of relevant values used in a valuation would also be a way of trying to resolve a blocked decision.

Widening Scope	increasing number of values considered relevant
Narrowing Scope	decreasing number of values considered relevant

In a state of value indifference in choosing a house to buy a deadlock breaking value like architectural quality could be introduced to try to obtain a dominant choice. Or in a state of value conflict a total ranking among values such as

Soundness > Cost > Amenity > Size Aptness > Accessibility > Utility

could be exploited to select the least important value which could be removed from the valuation. By eliminating the value *Utility* from the set of relevant values, a dominant alternative might emerge from the valuation. This process might be iterated by eliminating successively more important values until a dominant choice emerges. Plainly if such a process continues until only one value, the most important one is left, then either one alternative will emerge as the dominant one, or two or more alternatives will be left in a state of value indifference.

Each of these ways of resolving blocked choices has a certain appeal, but none seems to require to be tried first. For this reason using any one method for resolving blocked choices is potentially controversial. Use of another method might resolve a blocked choice a different way. Furthermore independent use of any of these methods might not be enough. They might be used in a number of possible combinations with different results. Plainly there are many ways of combining these methods, and it would seem likely that there are many other reasonable methods for resolving blocked choices which have not been considered. So there would seem to be wide scope for judgement. The Praxis model acknowledges this by imposing no general requirements on default rules for resolving blocked choices.

9. Conclusion

The Praxis model conceives of reasons for acting in terms of requirements, duties, proscriptions and values. It provides a framework for representing different kinds of reasons for and against alternatives. The first set specify what must be done. They are deduced from general requirements and the circumstances of the decision context. A second set specify what ought to

be done. They are deduced from general praxeological principles and often imply that conflicting courses of action ought to be done. These oughts arise out of duties, needs and interests of the reasoner. The two sets of what must and ought to be done generate the options which means end reasoning refines into alternative things to be done. A third set of considerations are proscriptions specifying what may not be done. They set constraints on practical decision-making by excluding alternatives which entail forbidden acts. A fourth set of considerations in the Praxis model are values. They provide the dimensions of assessment for alternatives which ought to be done. They succeed if they show that one alternative has more value than any others. If none of these considerations determine what is to be done all things considered, the Praxis model allows default rules for resolving blocked choices to be applied.

10. Acknowledgements

I would like to thank Mike Clarke, John Fox, Paul Krause and Saki Hajnal for discussions on these topics. Work on this paper was supported by a SERC project concerned with formalising reasoning for decision support systems. The grant holders were Mike Clarke of QMW and John Fox of ICRF.

References

1. L. Aqvist, *Introduction to Deontic Logic and the Theory of Normative Systems*, Bibliopolis, Naples, 1987.
2. I. Asimov, *Robot City*, Orbit, 1988.
3. D. Davidson, *The Logical Form of Action Sentences, in Essays on Actions and Events*, Clarendon Press, Oxford, 1980.
4. G. Harman, *Change in View*, pp, 23-35, MIT Press 1986.
5. D. Hume, *Treatise on Human Nature*, Clarendon Press, Oxford, 1978.
6. R.C. Jeffrey, *The Logic of Decision*, McGraw-Hill Book Company, New York, 1965.
7. D.V. Lindley, *Making Decisions*, John Wiley, London, 1985.
8. R.D. Luce and H Raiffa, *Games and Decisions*, New York, 1957.
9. R. Nozick, *Philosophical Explanations*, Oxford University Press, 1984.
10. N. Rescher, *Introduction to Value Theory*, Prentice Hall, 1969.
11. D.A.J. Richards, *A Theory of Reasons for Action*, 1972.
12. L.J.Savage, *The Foundations of Statistics*, Dover, New York, 1972.
13. G.H. von Wright, *Essay in Deontic Logic and General Theory of Action*, North Holland, 1968.