CHAPTER 15

Evaluation and Choice

Presupposes: the three-cell partition of mental features (Ch. 13, §2.3); the response-constitutive account of value (Ch. 14).

Establishes: the primary/added value distinction, individuated by the severing test; the structure of evaluation and choice; the constitution of the evaluative hierarchy; weighing as field-enrichment, and across evaluators as convergence under mutual registration; the account's convergence with rival theories and its answers to the standing challenges. All results are constructed; note 1 records the grades and the three empirical stakes.[1]

In brief. Why does a shirt your father once wore feel different from an identical shirt bought yesterday? The difference is in you — and this chapter supplies its mechanics. Total value is a composite: primary value (the direct pull of the thing itself) plus added value flowing backward from its source or forward from its effects. The instrument that keeps the distinction disciplined is the severing test: cut the association while holding the object fixed, and whatever dies was never the object's own — peel the label from the wine and the provenance-glow goes with it while the flavour's pull remains. The same mechanics answer the hardest question a response-based theory faces: when accurate responses conflict, there is no cosmic exchange rate waiting to be found. Deliberation is not measurement that fails for want of a unit; it is the widening of the field until it ceases to be tied — and between people, each party's taking in of the other's registration, with what such fields converge on the only fact a trade-off has.

---

1 Introduction

Why does a shirt your father once wore feel different from an identical shirt bought yesterday? Both are cloth; both serve the same function. The difference is in you — in the web of emotional associations that transforms a piece of fabric into something irreplaceable. The previous chapter established that emotion constitutes evaluative difference.

It did not explain how evaluation works in practice: how multiple responses with different intensities are resolved, how the result becomes a choice, and why people rank the same things so differently.

This chapter supplies the mechanics. Total value is a composite with a specific structure: primary value, the modality's direct response to the object, plus added value in two species — derived value flowing backward from the object's source, consequential value flowing forward from its effects.

The structure explains, within one mechanism, phenomena the existing taxonomies treat piecemeal: why identical objects can differ in value, why value changes without the object changing, why adaptive preferences have the profile they do, and, in the section this chapter adds to the sequence (§4.4), how deliberation escapes deadlock when accurate responses conflict. Beside what it supplies, the chapter returns no fact about the right trade-off between parties, only about what mutually informed fields settle on (§4.4), a limit every later chapter carries. The partition is not merely taxonomic. It is the mechanism by which the evaluative hierarchy is constituted and transformed, and every subsequent chapter presupposes it.

2 The Structure of Evaluation

2.1 From stimulus to evaluative response

Evaluation has a determinate structure, however rapidly it unfolds. A stimulus is detected by the sensory modalities. The detected distinctions are processed cognitively: the object identified, classified, and associated with stored representations, each association carrying its own possible emotional charge. The evaluative modality then responds to the resulting complex with attraction or repulsion. That response is the evaluative difference.

The process scales with the field. A simple evaluation of a shirt might involve liking its colour, the modality responding to a single sensory feature. A richer one adds the memory of your father wearing such a shirt: warmth, nostalgia, the pang of loss. Richer still, the

knowledge that it was produced by exploitative labour triggers a negative empathic response to the workers' suffering. The object has not changed; the evaluative field in which it is embedded has, so the overall evaluation is the resultant of all the oriented responses the field elicits.

Two features matter for everything that follows. The responses are constitutive rather than representational: the modality does not discover that the shirt is valuable and report its finding, but constitutes the value through its response. And the cognitive and evaluative operations are collaborative while distinct. A system with perfect cognitive processing and no evaluative modality would identify every association and register every fact without experiencing the shirt as attractive, repulsive, meaningful, or worthless.

2.2 Primary and secondary objects

The elements of an evaluative process divide in two. The primary object is what initiates the process, the shirt itself. The secondary objects are everything associated with it: the memory of your father, the brand's labour practices, the friend who recommended the shop. Each may carry its own evaluative weight, so what we call "the value of" a thing is rarely a single response. It is a composite, the resultant of the modality's oriented responses to the primary object and its secondary web. The distinction is structural rather than normative: in many cases the secondary objects dominate entirely, as when a plain envelope carrying a letter from a dying friend is evaluated almost wholly through what it carries.

3 Primary Value and Added Value

3.1 Primary (essential) value

The primary or essential value of an object is the emotional response elicited directly by the primary object, without association with any other elements — what the object is worth to you on its own, stripped of its web. Primary value is the baseline from which the hierarchy is built, and often the first response to register.

But immediacy is not primacy. Kahneman's (2011) "System 1" is defined by speed, not by which route the response travels: a fast response may be primary, the direct pull of a colour, or added, the automatic warmth an association discharges before any deliberation. Primary value names a route, System 1 a tempo, and the two axes cross — a point §6.1 uses.

3.2 Added (secondary) value and the severing test

Added value is any emotional response elicited by secondary objects, dividing into two kinds by direction. Derived value flows backward along the causal chain, from source to object: the shirt given by your father, the painting reattributed to Vermeer, the object unchanged while its derived value is transformed. Consequential value flows forward, from object to effects: the shirt that protects sunburnt skin, the canvas a billionaire will pay a fortune for.

The taxonomy needs an individuation criterion, because without one the primary/added line can be redrawn to suit any diagnosis. The Vermeer case invites the challenge directly: is the collector's surge of feeling added value on the painting, or the primary value of a different object, the-painting-as-a-Vermeer? If redescription can convert any added value into primary value, the distinction does no work.

The severing test. The criterion is counterfactual. A component of the total pull is added if it would not survive the severing of the association with the object held fixed, and primary if it would. Sever the attribution — same canvas, same pigments, the Vermeer identification withdrawn — and the surge collapses, so the component was derived. Peel the label from the wine and the provenance-glow goes with it while the flavour's own pull remains. The test is applied to components under intervention, which is what blocks the redescription manoeuvre: the objector may call the object what he pleases, but the intervention holds the object's intrinsic features fixed and cuts only the association, and what the cut kills was never the object's own.

The transformation caveat of §3.3 marks the test's one refinement, which must be stated with its discipline, or the caveat becomes an alibi.

The test is time-indexed: a component is added at a time if severing at that time kills it, primary at a time if it survives.

Transformation is accordingly a historical claim rather than a synchronic one. Where an association has retrained the primary response, severing no longer reverses the pull, so the test correctly reports a conversion of added into primary value. But to say so is to say the component was severable earlier and is not severable now, the conversion mediated by a traceable associative history. A transformation invoked with no history is inadmissible, since a component that was never severable is primary throughout rather than converted. This is what keeps the caveat from excusing every unwelcome verdict: it predicts a profile — severability decaying over exposure — distinguishable from primacy, which is severability's constant absence.

Stated as evidence, the admissibility conditions are three: a record of the exposure through which the conversion is claimed to have run; a dose profile, severability declining as the exposure accumulates; and the asymmetry the mechanism requires, conversion running forward with exposure and never reversing spontaneously in its absence. A transformation claim that can produce none of the three is rejected at the door.

So disciplined, the caveat is the account's theory of evaluative learning: the mechanism by which taste is formed and, run deliberately from outside, the mechanism marketing exploits (§6.1, §7.2). The time-indexing makes "primary" a stage-relative status, which the account accepts without strain. A pristine, history-free primary value would be the pre-social standpoint §6.6 disowns, and the framework's primary value was always the response's current direct route rather than its innocence.

3.3 The structural consequence

Total value is the composite of primary and added value, with one caveat: the relation is not always additive. Added value can transform the primary response rather than combining with it — learning that a

painting was made by a child prodigy may change the aesthetic response itself — so the formula is a simplification rather than an independence claim.

The consequence is striking. An object with negligible primary value can carry enormous total value: the painful procedure embraced for the life it saves, the worthless-looking object worth millions as a lost masterwork. The converse holds too, since the beautiful shirt made by child labour and the delicious food known to be poisonous are evaluated negatively overall. Added value overrides primary value and reverses rankings. This interaction makes the evaluative hierarchy dynamic, a constantly updated composite responsive to everything that enters the field; the dynamism is what makes deep persuasion possible (Chapter 21), since a newly enriched field constitutes a different composite.

3.4 Relation to the intrinsic/extrinsic value debate

The primary/added distinction intersects with the intrinsic/extrinsic distinction without being it. Korsgaard (1983) prised apart two senses of "intrinsic value", value as an end (final value) and non-relational value; M. J. Zimmerman (2001) developed the second as a monadic property supervening on intrinsic features alone, a reading Korsgaard's own cases of final value borne by relational features tell against.

Primary value is the nearest analogue of Zimmerman's intrinsic value, added value of the relational kinds, with two differences that favour the present cut. First, primary value is response-dependent, its variation across evaluators no more an objection than colour's variation across visual systems (Chapter 14, §2): the structural insight preserved, the ontological commitment replaced. Second, the primary/added taxonomy is compositional where intrinsic/instrumental is static. Total value is composed in real time as secondary objects enter the field, which is why the same object changes in total value without changing intrinsically, and why intrinsically identical objects differ. §3.5 sharpens the comparison into a prediction.

3.5 From taxonomy to partition: the derivation

As presented so far the distinction is a taxonomy, and a taxonomy invites the question convergence cannot answer: why must value divide this way? It can be derived as an exhaustive partition from results already established. Given the response-constitutive account (Chapter 14, §3.5), every component of an object's total value is an oriented response to it, and valence is never free-standing, being the orientation of a content (Chapter 13, §2.3).

The derivation. Ask of any evaluative response what content it rides. Either the content confines itself to what the encounter with the object itself delivers — the diffs registered in seeing, tasting, wearing, being with it — or it includes represented items the encounter does not deliver: source, history, effects, price, owner. Inclusion is binary; there is no third condition to be in. The first route constitutes primary value, which is why §3.1 could define it as the response "without association with any other elements", isolation being the encounter without importation. The second constitutes added value, its species falling out by the direction of the imported relation: backward yields derived, forward consequential.

One clarification guards the edge. Encounter-delivered character includes the relational diffs the encounter itself presents — the painting's contrast with its neighbours, the shirt's fit on this body — so the line runs between the delivered and the imported, cutting across the line between the intrinsic and the relational. The partition is exhaustive by its own logical form, which is what it rests on: any evaluative content either exceeds the object's registered character or does not. The diversity/equality partition divides on distinction against its absence, whereas this one does not descend from it; what the two share is the form of a contradictory pair, which is a resemblance in shape.

The derivation converts §3.4's positioning into a prediction. Final/instrumental classifies the *mode* of valuing; intrinsic/extrinsic the value's *ground*; primary/added the *route* of the response's content — three cuts at three levels, so the three cross. The collector's Vermeer and the father's shirt are valued finally, through the added route: extrinsic

final value, which the standard taxonomy accommodated only after Korsgaard prised its distinctions apart, arrives here as a predicted cell rather than a curiosity.

The partition also yields a structural diagnosis of the fitting-attitudes literature's hardest internal problem. In the demon case (Rabinowicz and Rønnow-Rasmussen, 2004), there is excellent reason to admire the saucer of mud, yet the mud is not thereby valuable, so the analysis owes an account of why the demon's reasons are of the "wrong kind".

The route partition supplies one. The demon's reasons enter on the consequential channel, since their content represents the effects of admiring rather than the mud's encounter-delivered character, while admiration is a response constituted by the primary route. The demon's reasons therefore ride a route that cannot constitute the response they demand. They can make admiring useful; they cannot make it registration. Wrong-kind reasons are wrong-route reasons.

This is the diagnosis of Chapter 14 (§6.4) at a finer grain rather than a second one competing with it. There the demon's reason was placed by axis, bearing on the value of hosting the response rather than on the accuracy of its content. Here the same reason is placed by route, entering on the consequential channel of a response the primary route constitutes. The axis analysis says what the demon's reason fails to touch; the route analysis says by which channel it arrives. Both are needed, and neither is prior.

A challenger to the partition has two options. She must reject response-constitution itself, or exhibit an evaluative response whose content neither confines itself to the encounter nor imports anything beyond it — composite objects included, since a configuration the encounter presents is encounter-delivered content (§7.6). The two-axes architecture leaves nothing for such a content to be. With the partition derived, the mechanics of choice can be set out.

4 The Structure of Choice

4.1 Choice as acknowledgement of preference

An act of evaluation leads to a choice: the acknowledgement of a preference between options, requiring at least two evaluative responses of differing intensities. Between two shirts, we select the one eliciting the stronger positive response, or, in an all-negative scenario, the weaker negative one. This inverts the standard theory of rational choice (von Neumann and Morgenstern, 1944), on which preferences are primitive inputs. Here preferences are constituted, outputs of the evaluative process, so the theory that models their aggregation is supplied with an account of where they come from.

Chang's (2002) parity cases receive a mechanism on the same basis. Two options are on a par when their total values are close in intensity while constituted by responses directed at different dimensions, so no single comparison resolves the ranking. Parity is real as a description of the pre-choice field. Under compelled selection the modality resolves it, because the agent must orient to act — resolved by choosing, never eliminated by it.

Two commitments need their epistemic weight stated, or the account risks truth by bookkeeping. First, the principle that choice follows the strongest total value would be a tautology if total value were read off the choice. It has content because total value is defined by its constituents, primary response and added associations, each identifiable in principle independently of the choice: by report, by physiological measure, by response to the components in isolation. So defined, the principle is an empirical claim that would fail wherever independently measured total value and compelled choice came apart systematically — though its content is modest, since the composition rule assigns no weights to the constituents, and until it does the principle predicts only that choice tracks the constituents' sum in sign and rough magnitude, not that it tracks a determinate function of them.

The shape of a test is concrete. Take a pair of options differing in a single added association — an owned mug against an identical unowned

one, the endowment paradigm — and measure the components separately: the valuation of the bare good, then the ownership premium elicited independently. The principle predicts that compelled choice tracks the summed measures; the stable premiums of the endowment literature are the nearest existing data of this form (Kahneman, 2011).

Second, the forced-choice diagnostic (§5.2) claims that compelled selection under stated conditions — repeated choice, relevant information present, situational constraint controlled — reveals a stable ordering. The construction-of-preference literature (Lichtenstein & Slovic, 2006) documents reversals tracking elicitation format, obtained outside those conditions. The account predicts they attenuate as the conditions are approached. Evidence that reversals shrink with repetition and market experience exists (Plott, 1996; List, 2003) and is consistent with the prediction without yet testing it under the diagnostic's own conditions, so the prediction is carried as a stake: persistent, undiminished reversal under those conditions would show that nothing is being revealed.

The diagnostic is staked on that convergence, with one confound to control. Having chosen, the agent acquires added value for the option taken, the derived worth of its being hers, which stabilises later choices and can counterfeit a pre-existing order (Chen & Risen, 2010). Convergence must therefore be read across choices insulated from that feedback. A single forced choice remains ambiguous between disclosing an ordering and constituting one; only insulated repetition discriminates.

4.2 Three kinds of choice

Choices classify by the evaluative profile of the options. All-positive: both attract, one more strongly — two songs you enjoy. Mixed: one attracts and one repels, as with the well-paid, soul-destroying job against the poorly paid vocation; this is the most common form in moral life, and typically the most complex. All-negative: both repel, and you choose the lesser evil — the person atop a burning building, choosing between jumping and staying, moved by the difference in intensity of two fears rather than by attraction.

4.3 Reflex, instinct, and intention

Not all behaviour involves this process. Reflexes bypass thought and emotion entirely, since the knee's kick travels a peripheral circuit; no evaluation occurs, so reflexive behaviour falls outside the evaluative domain. Instinctive actions involve emotion without the comparison of alternatives, as the scream at the attacking dog follows the response directly, evaluative but not deliberate. Intentional actions involve the full process: options represented, each evaluated, the responses compared, the strongest acted on. The triad adds a finer grain than Kahneman's binary, whose System 1 spans both reflex and instinct, where the difference between them matters for moral responsibility (Chapter 18).

4.4 Deliberative deadlock and field-enrichment

When two accurate registrations pull opposite ways, there is no scale to read off which is heavier. This section says what deliberation does instead — it widens the field until the field tips — and what is left when it will not tip.

The account must now meet the objection it has seemed to invite since Chapter 14: that a theory on which value is constituted by responses possesses a fact-maker and no scale. Where registrations conflict, as with five patients' suffering against one patient's death, or a friend's present hurt against his future ruin, perfectly accurate empathy delivers opposed pulls of comparable force, so the demand arrives for the correct exchange rate: the fact about which outcome is worse, independent of anyone's weighing.

Why the demand dissolves. An exchange rate valid independently of every weigher is a response-independent evaluative fact, the vacant clause of Chapter 14 (§3.6) in its practical costume, which dissolves as its predecessors dissolved. There is no cosmic scale for the same reason there is no free-standing badness: a ranking no possible evaluative system produces is a ranking composed of nothing. The theorem reaches procedures as well as outcomes, since a procedure every

possible party must register as fair would be the same vacant item one storey up.

The asymmetry this leaves should be stated, because the account leans on it in both directions. Within one evaluator the components of a composite do resolve: they are orientations of a single register, compelled selection is the operation that resolves them, while the constituents are separately measurable (§4.1). Across evaluators no such operation exists, since valence is proprietary to the register that undergoes it (Chapter 14, §3.5) and no register undergoes two parties' pulls at once. The demand for an exchange rate is therefore a demand for a resolution with nothing to perform it.

The demand's strongest form arrives as a consistency argument, which deserves separate audit. Mere addition (Parfit, 1984) and the spectrum arguments descended from it proceed by transitivity and dominance: if each option in a chain is no worse than its neighbour, the endpoints stand ordered, and a theory that refuses the ordering is charged with incoherence.

The reply distinguishes the constraints from what they are asked to do. Transitivity and dominance are orthotic constraints on a comparison scale already in place: once a practice of ranking is adopted, violating them is a defect of the practice, which the account keeps in full force. What they cannot do is institute the scale, because each pairwise step of the spectrum presupposes the commensurability of its adjacent cases, which is the point in dispute. The argument's chain is built from the very comparisons whose availability it was meant to prove.

Where the parties' registrations sustain the pairwise comparisons, the constraints bind inside the practice those registrations constitute. Where the registrations give out, the chain has no links, so the refusal of the endpoint ordering is principled: proprietary registrations do not sum across bearers (Chapter 14, §6.5). The cost of saying so is the one this section has already owned — a field that remains level after full enrichment leaves a choice among tragedies, with the residual work belonging to procedure.

What the demand misdescribes, this chapter's mechanics already describe. A deadlocked choice is not a question awaiting a missing answer. It is a state of the evaluative field, opposed responses in equilibrium, and fields exit equilibrium by the only route they ever had: enrichment. New secondary objects enter — what the choice does to the surviving families, to trust in the institution, to the precedent governing the next hundred cases, to safety, to liberty — each carrying added value of the kinds §3 taxonomised, until the composite tips and motion resumes. Deliberation is not measurement that fails for want of a unit. It is the widening of the field until it ceases to be tied.

The practice confirms the analysis at its terminus. Where enrichment runs out and the options remain level, institutions resort to the lottery, and a lottery is intelligible only on the present account, as one more added value rather than an exemption from the theory. Procedural fairness is a secondary object the parties' registrations can converge on when nothing else differentiates the options, and whether they converge is itself contingent, constituted by their actual responses like everything else in the field.

The practice shows both outcomes, and both are data for the account. Some parties register the coin as the one impartiality left standing. Others register a lottery over lives as an abdication, a refusal to weigh experienced as an insult to the stakes, which is why triage boards and courts have as often rejected sortition as adopted it. A theory of correct weights can explain neither party: not the ones who accept a coin, since beings who possessed the weights would never need one, nor the ones who refuse it, since their refusal registers nothing on that theory but error. On this account both are what they seem — registrations, contestable, belonging to the procedures that follow.

The legitimacy criterion. Enrichment requires a discipline, or the account licenses the manipulator, whose whisper also adds value and also tips fields. The discipline is the standard fixed in Chapter 14 (§2.2): an added value is legitimate insofar as it bottoms out in real registrations. "What this does to the community", pursued downward, terminates in feeling creatures — the families, the future patients, the citizens whose trust is at stake — and is empathy cast over a wider

causal field. The demagogue's installed association is a harder case than that sentence suggests, since a demagogue's whisper usually does bottom out in real registrations — real fear, real loss — and what makes it manipulation is a misreading: the association is keyed to a picture of the affected parties that is false or incomplete, so the criterion must be read as Chapter 21 (§7.6) reads it, as the fittingness of the added value to the registrations it invokes. So read, the whisper that terminates in no affected party and the whisper that terminates in a misdescribed one are classified together, as corruption, by the accuracy-vocabulary that classified the mannequin case.

Weighing, so understood, inherits the account's full error-theory. One weighs wrongly by excluding registrations from the field, or by reading falsely the ones admitted, which is historically how the great moral failures proceeded (Chapter 17, §7.1) — never by miscomputing a sum that was there to compute. Where a fully enriched, fully accurate field remains in equilibrium, the account says what no rival theory says: there is no further fact within that field; the residual work belongs to the operation across fields that the next paragraph names, the deep persuasion of Chapter 21 and the lawmaking machinery of Chapter 22, which are this mechanism at social scale.

Enrichment across evaluators. Across evaluators, enrichment has a form of its own, and the political sequence runs on it. Picture two registrations in conflict as two screens in one hall: the child's suffering on one, the five patients' on the other. There is no third screen on which the two appear weighed against each other. The weighing is in the viewers.

Call a field mutually informed when each party accurately registers the other's registration — when the surgeon weighing five against one has taken in what each of the six undergoes, and each of the six, so far as they can, what the others do. The other party's pull enters this party's field as an empathic secondary object (Chapter 14, §4.2), which is the one object the demand for an exchange rate was looking for in the wrong place.

A verdict on a trade-off is then whatever mutually informed fields settle on. No scale is applied from outside; the settling is reached by deep persuasion, and found nowhere else. The demand for a fact was, underneath, a demand for something shareable, and a shared response supplies that as well as a shared object does. The account stakes one claim on this: mutually informed fields converge more than they diverge. It is the thinnest empirical bet the account makes, marked here as well as in the Register. The contact and deep-canvassing studies Chapter 21 draws on already test the claim, within their limits.

The mechanism has an ancestor. Smith (1759/1976) built moral judgement from mutual sympathy, each party imaginatively taking up the other's feeling and the two adjusting until their sentiments concord, and then idealised the concord into a single impartial spectator whose verdict is the standard. Mutually informed fields are Smith's mechanism without the idealisation: the concord is reached in the parties, with no constructed spectator to reach it for them; where it is not reached, the account records a level field, with no verdict a spectator would have supplied. Harsanyi's (1955) impartial observer, who chooses as though he might occupy any party's position with that party's preferences, is the same move in decision-theoretic dress; the account takes the imagining and declines the sum Harsanyi builds on it, for the reason Chapter 14 (§6.5) gives.

Three outcomes are possible. Where the fields converge on what to do, the convergence is the fact, and nothing further is wanted. Where they do not, they more often converge one level up, on how to decide: the triage rule, the lottery, the vote. Scanlon's individualist restriction, which tests a principle against each person's complaint and never against the sum (1998; Chapter 16, §4.2), is a convergence of this kind stated as a rule. That is the convergence on procedural fairness described above, a registration in its own right.

Where neither level converges — a fully informed field level on the outcome and level on the procedure — the account's residue begins. There is no further fact, only a choice among tragedies, and what remains are Chapter 21's other two modes, enforcement and management. A theory of correct weights would settle none of these

cases either. It would assert that they are settled; the screen in the empty hall shows what that assertion comes to.

Where the argument stands: evaluation has been given its structure (§2), total value has been partitioned into primary and added and the partition derived (§3), and choice has been analysed as compelled selection with deadlock exited by enrichment (§4). What those responses build, across a life, is the evaluative hierarchy.

5 The Evaluative Hierarchy

5.1 Formation

Differing emotional intensities toward different objects generate a hierarchy of value: an ordered ranking of objects by the intensity and orientation of the responses they elicit. Each person's hierarchy is unique, shaped by developmental history, accumulated association, together with the involuntary delivery of emotional content through subconscious processes she does not control. None is static. A trusted friend's betrayal destroys derived value and the friend's position shifts; a career valued for prestige reverses when it is found to harm. The dynamism follows from the composite structure, since any change in the secondary objects can transform the ranking.

5.2 The forced-choice diagnostic

Imagine five faulty devices, each containing a different object — a plant, a snake, a dog, a child, and a precious shirt — each about to activate and destroy its contents. You must decide in which order to rescue them. Different people choose differently, because the choice depends on each object's total value in the rescuer's hierarchy. Primary value is the empathic response the object triggers directly, strongest for the child and weakest for the plant in most human hierarchies. Added value enters where a billionaire's fortune for the shirt may lift it above the plant, or the dog — would anyone place it above the child? The scenario is an intuition pump and is offered as one; what makes the diagnostic a diagnostic is the set of evidential conditions §4.1 fixed —

insulated repetition, controlled feedback — under which a compelled choice discloses an ordering already in place.

Now replace the objects with humans and a dog, then introduce secondary factors. The man is your brother (derived); the woman knows the cure for cancer (consequential); you hold a principle that children are saved first, an orthotic assessment interacting with the empathic evaluation. The hierarchy shifts and complexifies with each importation, so the ranking that emerges is constituted by the totality of responses, each with its own intensity.

The philosophical yield is specific. The hierarchy is a lived structure that manifests in action and cannot be fully articulated in advance of what this chapter calls the forced-choice diagnostic (FCD): the revelation of hierarchical position under the pressure of compelled selection. The evaluator may not know which object she values most until she must choose. The hierarchy is constituted by evaluative responses rather than by cognitive judgements, so the responses emerge under the situation's pressure, sometimes surprising the evaluator herself. The FCD's evidential conditions and its constitutive limit were fixed at §4.1; subsequent chapters reference it by this label. (Gibbard's (2003) plan-expressivism marks the connection to normative theory: if normative judgements are plan-endorsements, the plans are products of the hierarchy the FCD reveals, which is the foundation of normative life.)

5.3 The subconscious hierarchy

A parallel hierarchy operates beneath awareness. Emotional patterns stored in the subconscious (associations with memories, fears, desires, traumas) carry their own evaluative weight and continue to shape reactions, preferences, and choices whether or not they ever surface (Bargh & Chartrand, 1999). A person who inexplicably avoids a kind of social situation may be responding to a buried derived-value connection between that situation and a childhood humiliation. Psychological analysis aims to uncover these layers; the subconscious hierarchy is the involuntarily constituted base on which the conscious one is built.

6 Engagement with Rival Accounts

6.1 Kahneman and dual-process theory

System 1 (fast, automatic) and System 2 (slow, deliberate) map onto the account with significant convergence, to which the account adds two things the framework lacks.

First, it separates two axes the fast/slow contrast runs together: tempo and route (§3.1), which cross. A fast response may be primary, the immediate pull of a colour, or added, the halo an association discharges automatically. A slow one may be primary, attentive savouring of a wine's own character, or added, the deliberate weighing of provenance and consequence. Kahneman's System 1 is therefore not one evaluative kind but two, distinguished by route beneath a shared tempo — the diagnostic the value taxonomy supplies where fast/slow cannot. When System 1 says "attractive" and System 2 says "dangerous", the conflict may be primary-versus-added or added-versus-added, which call for different corrections. Marketing does its work here, installing fast added value that wears the immediacy of a primary response (§7.2).

Second, the analysis identifies what both systems share. Both operate through the evaluative modality, since System 2's reasons are relevant only because they engage it: a reason that triggers no evaluative response is inert, which in the model of Chapter 14 (§A.5) is a field with no body in it. The difference between the systems is breadth of field, not kind.

6.2 Damasio and the somatic marker hypothesis

Damasio's (1994) somatic markers — bodily emotional signals that mark options as desirable before deliberate reasoning begins — converge substantially with the account, which adds the taxonomy his markers lack. A marker says only "this feels good". The account differentiates whether the feeling is constituted by primary, derived, or consequential value, a differentiation that matters clinically: a gambler's marker may respond to the thrill of the bet, which is primary, while

failing to register the financial consequences, which are consequential. The marker is a composite, and a composite can be analysed.

6.3 The rationalist tradition

Decision theory from expected utility (von Neumann and Morgenstern, 1944) to its refinements (Broome, 1991) treats choice as calculation, with emotion as an input given from outside. The account inverts the relationship. The calculation is an orthotic operation on evaluative content the modality has already constituted; without the oriented responses there would be nothing to calculate, no option better or worse than any other.

Broome's consistency requirement illustrates the point. Consistency among preferences is a real orthotic achievement, yet it says nothing about where preferences come from, so a perfectly consistent hierarchy built on distorted responses — a phobic ordering in which the harmless is ranked as threatening — is rational by Broome's criterion and defective by this account's, which is the criterion doing the work when we call it defective.

6.4 Scanlon and the buck-passing account

Scanlon (1998) holds that to be valuable is to have properties that provide reasons — value passes the buck to its grounds. The account agrees that value is no monolithic property, and locates the divergence with a thought experiment: in a world of beings with perfect cognitive faculties and no evaluative modality, do the properties of objects provide reasons? On Scanlon's account they do, waiting to be recognised; on this analysis the world contains no reasons at all, because the modality that constitutes the differences in virtue of which anything could count as a reason is absent. The buck-passing account correctly identifies the features that engage the modality and misidentifies them as self-sufficient sources: they are sources of reasons only because a modality exists that responds to them.

6.5 Bratman and planning agency

Bratman (1987) shows that intentions are not momentary desires but commitments that structure future deliberation; the account supplies what he leaves open, the sustaining mechanism. An intention is a sustained position in the evaluative hierarchy, maintained by the modality's continued positive response to the prospect of completion and negative response to abandonment, losing its structuring power when that orientation dissolves. This explains the familiar discovery that one no longer intends what one thought one intended: the cognitive commitment persists while its evaluative grounding has gone.

6.6 Nussbaum, Elster, and adaptive preferences

Nussbaum (2000) raises the challenge the account must confront directly. If hierarchies are constituted by responses shaped by conditioning, how are real preferences distinguished from adaptive ones, preferences shaped by oppressive conditions to endorse those conditions?

The primary/added distinction supplies the profile. The woman who prefers subordination does not find subordination intrinsically attractive; her preference is constituted by added value — the consequences of non-compliance (punishment, ostracism, insecurity) and the derived associations of the role (approval, belonging). Elster's (1983) sour grapes exhibits the same mechanism, since the grapes' primary pull was never altered, only outweighed by the consequential value of repeated failure. Where the primary response remains intact and merely outweighed, the diagnosis is operationalisable by the removal test: withdraw the externally imposed secondary objects — the punishment, the ostracism, the threatened condemnation — and a preference sustained by added value alone collapses, as the fox's indifference would with the grapes in reach.

The harder cases do not fit the suppression model, and §3.3 has already said why: added value can transform the primary response rather than override it. A person socialised from birth into a subordinate role may have no intact primary response waiting to re-emerge, since

the conditioning has shaped what her responses are — a transformation in the time-indexed sense of §3.2, carrying its history rather than pleading it. For such cases the removal test misfires, because nothing contrary re-emerges, and its error lies in the tacit assumption of a pristine, pre-social primary value constituting the person's real standpoint. The response-constitutive account recognises no such standpoint (Chapter 14, §6.4): there is no vantage outside every evaluative formation from which a person's settled responses can be declared not truly hers.

Two resources support the critique. What grounds the critique of a transformed adaptive preference is therefore not deviation from a natural baseline. It is the two resources that support evaluative criticism throughout the account.

The first is intra-agent consistency. The adapted preference characteristically fails to cohere with the agent's wider commitments, with the empathic response she extends to others in her position, and with the response she forms when the full range of options is before her. The second is convergence, since the direction of her conditioning is one against which human responses are otherwise near-universal. Neither resource pronounces her incorrect against a stance-independent fact, and neither needs to; both license the empathic verdict that her condition is bad and the practical project of altering it.

Deep persuasion is that project's mechanism, now described accurately. It does not restore a buried primary value. It enriches the field (§4.4), introducing new associations and making foreclosed options available, so that a different composite can form — one that coheres with the commitments from which the conditioning had isolated her.

7 Objections and Replies

7.1 The rationality objection

Objection. The account reduces choice to emotional response, leaving no room for rational deliberation.

Reply. It identifies what rational deliberation operates on. Reasons are considerations the cognitive system identifies as relevant, and a reason moves an agent only if it engages the modality — a reason without evaluative engagement is an inert proposition (§6.3, §6.4). Deliberation is the cognitive processing of evaluative content, not the source of it — and, in the deadlocked case, it is the enrichment of the field (§4.4), which is cognition and modality collaborating, not cognition replacing feeling.

7.2 The added value manipulation objection

Objection. If total value is a composite, marketing is value-manufacture: celebrity association, status signalling, and manufactured scarcity inflate a product's added value without touching its primary value.

Reply. The objection identifies a real phenomenon; the account supplies both its diagnosis and its discipline. The inflated value is real as felt, since the consumer really is more attracted, while it is located in the secondary objects, which the severing test exposes: peel the association and watch what dies.

The normative line is the legitimacy criterion of §4.4: added value is legitimate insofar as it bottoms out in real registrations — bottoms out in, not merely produces. The manipulator's whisper produces a registration in the person he moves and terminates in none in the party it invokes, which is the difference. The marketer's installed association, pursued downward, terminates in no affected party's condition, which classifies it with the demagogue's whisper rather than with the surgeon's warning. A consumer equipped with the taxonomy can audit her own composite; a public equipped with it can audit its persuaders (Chapters 21–22).

7.3 The commensurability objection

Objection. The hierarchy assumes all objects rank on a single scale of intensity, but some values seem incommensurable — a friendship and a career cannot be meaningfully compared.

Reply. The hierarchy requires no common metric available to reflection. It requires only that, when a choice is forced, the modality orients — which it does, whether or not the agent can articulate a scale. Unrankable options are Chang's parity, placed at §4.1: where the responses are matched there is no antecedent ranking for the forced choice to disclose, so the choice constitutes a resolution rather than reporting one. The deliberative version of the problem, how conflicting registrations are weighed, is answered by enrichment (§4.4): deadlock is exited by widening the field rather than by discovering the missing unit, and where a fully enriched field remains level there is no further fact to find.

Chang's own later development of the parity cases deserves the comparison, because it offers a rival mechanism for what happens when enrichment has done its work and the field is still level. On her account (Chang, 2017), hard choices are occasions for normative self-creation: the agent can put her will behind one option and thereby confer on herself a reason she did not have before — commitment as the source of new normativity rather than the discovery of old.

The present account keeps the phenomenon and relocates the mechanism. Committing is itself an entry into the field. The act places added value on the chosen side, identity-derived because the option becomes hers, and consequential because abandonment now carries costs that drifting never did; that added value is why commitment binds where velleity does not. What Chang reads as the will manufacturing a reason, the composite reads as the will supplying one more secondary object whose added value tips and then stabilises the total.

The difference matters downstream. Chapter 21's deep persuasion inherits enrichment as its mechanism, and a will that could mint reasons by fiat would reopen the door the manipulation discipline (§4.4) closes. Added value deposited by commitment bottoms out in real registrations, the agent's own, and passes the same audit as every other entry.

7.4 The cultural variation objection

Objection. Hierarchies vary enormously across cultures; if they are constituted by responses, value is completely relativised.

Reply. The account claims a common structure, not a common hierarchy: all hierarchies are constituted by oriented response to primary and secondary objects, differentiated by route into primary, derived, and consequential value. Content varies; structure does not; and the taxonomy is a tool for analysing any hierarchy, not a prescription of what should be valued. Where hierarchies collide morally, the anchoring of Chapter 14 (§6.4) and the machinery of §4.4 apply.

7.5 The reductionism objection

Objection. The account reduces value to emotion, ignoring moral principles, rational commitments, and long-term projects.

Reply. Principles, commitments, and projects contribute to the hierarchy through the modality. “Never harm a child” has force because the modality responds with strong negative orientation to the prospect; “I will finish my degree” because it responds positively to completion. These are forms of added value operating within the hierarchy, whose mechanism the account identifies through which they acquire the importance the objection correctly attributes to them.

7.6 The organic unities objection

Objection. G. E. Moore (1903) held that the value of a whole must never be assumed to equal the sum of its parts, since the consciousness of a beautiful object exceeds the value of the beauty plus the value of the bare consciousness. If organic unities are real, the compositional formula of §3.3 is false in principle, taking the severing test with it: cutting an association changes the whole, so the “components” were never constituents at all.

Reply. The account agrees with Moore about the phenomenon and departs over its mechanism. §3.3 already flags the additive formula as a simplification, since added value can transform rather than combine. What the framework adds is an explanation Moore left brute: the bearer

of value is the oriented response, so the modality can respond to a configuration as such. A composite content can elicit a pull no component elicits in isolation, because the response's object is the Gestalt, which is a content in its own right (Chapter 13, §2.2, on composition as a categorical transition). Organic unities are therefore predicted and mechanised where Moore could only assert them.

The severing test survives because it never presupposed additivity. It is an intervention asking a classificatory question — does this component of the pull survive the cut? — and where the interaction is organic it returns a verdict finer than the case's surface suggests, with the fineness staying inside the partition of §3.5 rather than adding a cell to it.

Configurational value is primary value whose object is the configuration. The Gestalt, as the reply has already said, is a content in its own right, and its contrast with its components is an encounter-delivered relational diff of the kind §3.5's edge-clarification admits. The pull that belongs to the arrangement and to neither relatum alone therefore rides the primary route, remarkable for its composite object rather than for a third channel.

That verdict identifies the cases in which value must be assessed at the level of the whole. It also explains the transformation cases of §3.3 as configurations in which an imported association has been absorbed into the object's registered character, absorption being what conversion means on the time-indexed test (§3.2): the once-imported is now encounter-delivered, and the verdict "primary on the composite" records the completed conversion. What is declined is only Moore's interpretation, since no non-natural property of the whole needs positing.

8 Conclusion

The primary/added distinction is the chapter's most consequential result, because it is the most widely applicable. The great moral conflation operates through it (Chapter 16), the incompleteness of free will is explained by it (Chapter 18), while the possibility of deep

persuasion is grounded in it (Chapter 21). The evaluative hierarchy is the medium in which all of these occur.

The chapter yields four results, all constructed.

The primary/added partition. Total value is the composite of primary value and added value in its two species, individuated by the severing test and derived at §3.5 as an exhaustive partition by logical form. A challenger must reject response-constitution or exhibit a content that neither confines itself to the encounter nor imports beyond it.

The structure of choice. Preferences are constituted, not primitive; choice acknowledges the strongest total response; the forced-choice diagnostic reveals the hierarchy under stated, falsifiable conditions and constitutes resolutions where responses are matched.

The evaluative hierarchy. Unique, dynamic, layered above a subconscious base the agent does not control.

Weighing as field-enrichment. Deliberative deadlock is equilibrium, exited by the entry of new secondary objects. Added value is legitimate insofar as it bottoms out in real registrations, while the lottery is one more added value, its acceptability constituted by the parties' registrations rather than guaranteed by fairness. Across evaluators, enrichment takes the form of mutual registration, so a verdict on a trade-off is what mutually informed fields settle on. Where they converge on the outcome, the convergence is the fact; where only on the procedure, the procedure is; and where on neither, there is no further fact — the residue named at §4.4, where Chapter 21's other modes remain.

If evaluation is structured in this way, with the hierarchy constituted by responses the agent does not control, a question presses. Have the ethical traditions correctly identified the evaluative operation that constitutes morality, or have they systematically confused the tool with the thing it measures? The next chapter argues that they have, and names the confusion.

Grades and Dependencies

1. The three empirical stakes. All four results are *constructed*, on the response-constitutive account (Ch. 14) and the two-axes architecture (Ch. 13). The severing test individuates the partition: a component is added at a time iff severing the association, object held fixed, kills it — with the three admissibility conditions on transformation claims (§3.2). Three commitments are staked empirically: choice tracking independently measured total value (endowment-paradigm form) and FCD convergence under insulated repetition (§4.1) — persistent reversal under the diagnostic's own conditions would falsify — and the convergence of mutually informed fields (§4.4), staked on the contact and deep-canvassing evidence of Chapter 21 (§6.2). The Premise Register carries the chain.

References

Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54(7), 462–479.

Bratman, M. E. (1987). *Intention, plans, and practical reason*. Harvard University Press.

Broome, J. (1991). *Weighing goods: Equality, uncertainty and time*. Blackwell.

Chang, R. (2002). *Making comparisons count*. Routledge.

Chang, R. (2017). Hard choices. *Journal of the American Philosophical Association*, 3(1), 1–21.

Chen, M. K., & Risen, J. L. (2010). How choice affects and reflects preferences: Revisiting the free-choice paradigm. *Journal of Personality and Social Psychology*, 99(4), 573–594.

Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. Putnam.

Elster, J. (1983). *Sour grapes: Studies in the subversion of rationality*. Cambridge University Press.

Gibbard, A. (2003). *Thinking how to live*. Harvard University Press.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4), 309–321.

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Korsgaard, C. M. (1983). Two distinctions in goodness. *Philosophical Review*, 92(2), 169–195.

Lichtenstein, S., & Slovic, P. (Eds.). (2006). *The construction of preference*. Cambridge University Press.

List, J. A. (2003). Does market experience eliminate market anomalies? *Quarterly Journal of Economics, 118*(1), 41–71.

Moore, G. E. (1903). *Principia ethica*. Cambridge University Press.

Nussbaum, M. C. (2000). *Women and human development: The capabilities approach*. Cambridge University Press.

Parfit, D. (1984). *Reasons and persons*. Clarendon Press.

Plott, C. R. (1996). Rational individual behaviour in markets and social choice processes: The discovered preference hypothesis. In K. J. Arrow, E. Colombatto, M. Perlman, & C. Schmidt (Eds.), *The rational foundations of economic behaviour* (pp. 225–250). Macmillan.

Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics*, 114(3), 391–423.

Scanlon, T. M. (1998). *What we owe to each other*. Harvard University Press.

Smith, A. (1759/1976). *The theory of moral sentiments* (D. D. Raphael & A. L. Macfie, Eds.). Clarendon Press.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.

Zimmerman, M. J. (2001). *The nature of intrinsic value*. Rowman & Littlefield.