Moral Philosophy:
An Abstract Approach

Document # PLPC-150020
Version 0.1
October 14, 2009

Available on-line at:
http://andrew.hammoude.1.byname.net/PLPC/150020

Andrew Hammoude
Contact: http://andrew.hammoude.1.byname.net/ContactMe
Web: http://andrew.hammoude.1.byname.net

Contents

1 Introduction

This is a series of essays about Moral Philosophy. In this series we have as our ultimate goal an ambitious objective: the formulation of a Unified Theory of Morality. We are seeking a theory that in principle answers *all* moral questions, in the same way that unified theories in other domains answer all questions within their domain.

If we can succeed in this, we will then be equipped to resolve *any* question or dilemma that we would commonly categorize as a moral question, or moral dilemma. We will be able to resolve intense and highly conflicted debates about such things as abortion, euthanasia, capital punishment, same-sex marriage—about essentially anything we consider to have moral content. We will have the intellectual framework and methodology to analyse and resolve such questions definitively.

In order to accomplish this there are some perennial philosophical questions we must address. Among the questions we must answer are: What is the meaning of morality? What is the meaning of right and wrong? What is the meaning of good and evil? Do these terms in fact have any meaning at all? And supposing they do, is there any sort of absolute or universal morality? And if so, what might that be?

The development will take the form of a series of linked essays, each building on the ones before, and each presenting one or more new ideas. It is the summation of these ideas that will give us something new and interesting.

1.1 Work-in-progress

This is work-in-progress. The intended essays are in various stages of construction—some in draft form, some consisting only of placeholder notes, and others at concept level only. At this point only the first essay exists, and only in draft form. This first essay is titled, "Defining an Abstract Morality," and appears below.

However an initial working Roadmap is available, providing an overview of the major anticipated essays, with placeholder notes to describe the general topic of each essay. For the interested reader, this will provide a general sense of where we are going. The Roadmap is available at:
http://andrew.hammoude.1.byname.net/MoralPhilosophy

This work represents the ideas and thinking of one man, and none of us has all the answers. To bring this work to maturity will require ideas, commentary and criticism from others. Everything written here is open to challenge. If you have thoughts to share, please let me know.

1.2 Audience

An issue an author must address in writing something of this nature is defining the intended audience. People's intellectual capability has multi-order-of-magnitude variability, from those who really don't know how to think at all, to those who process abstract ideas as easily as breathing. One could write a sophisticated presentation targetted to this latter audience, but then the target audience is quite limited.

Or one can write for a more general audience, but then the intellectually sophisticated reader will likely be exasperated by the excessive amount of (for him) superfluous exposition.

My decision is to write for a general audience. And this is for an important reason: my goal is to influence our cultural ideas and thinking about morality. And beyond that, perhaps even to influence *the way we behave*. And this possible influence on behaviour is in no way limited to an intellectual elite. On the contrary, it has relevance to all persons, everywhere. And so for the reader whose mind at each turn leaps forward like a greyhound, I ask for your patience.

2 Defining an Abstract Morality

We begin by defining something we will call an *abstract morality*, or when no confusion can arise, just a *morality*. This is an artificial thing we are making up. It bears some resemblance to the real world; it's defined using terms that look the same as terms used in real life, and its structure somewhat resembles our notion of what morality might mean in reality, but it's not real. It's an abstraction.

Formally, we define an **abstract morality** as a mapping from the *set of actions* into the two-element abstract set $\{\textit{right}, \textit{wrong}\}$.

What does this mean? What is the "set of actions"? What is a two-element abstract set? And what is a mapping? We'll clarify each in turn.

The set of actions

By the "set of actions" we mean all actions within a particular class of action. Roughly speaking, we mean *any action taken by an agent, that may impinge in some way for good or ill against an external entity*.

We are using the word "action" in the ordinary sense of the word, meaning the exertion of some sort of capability to cause something to happen, or be done or effected.

By "agent" we mean the perpetrator of the action. The agent may be a single individual, or a group acting in concert. The external entity also may be a single individual, or a multiplicity of individuals. But whether an individual or group agent, for technical reasons we impose the requirement that the agent have some cognitive awareness of the impinging effect of the action against the external entity. We will also impose the technical requirement that the external entity be sentient, i.e. capable of subjective experience.

So an example might be an individual person committing armed robbery against another person, or a national government sending its armed forces to war against a foreign nation. The former is a conceptually simple action with a single agent and a single affected entity; the latter a vastly more complicated group action, with affected entities potentially numbering in the millions.

But though these are widely differing examples, they both meet the criteria for inclusion in our set of actions. In each case the action consists of the exertion of a capability held by the agent, and in each case the action impinges against one or more external entities. In each case the affected entities can be expected to experience the action in subjective terms, and in each case the agent can be expected to have some conceptual awareness of this.

The set of actions consists of *all* possible actions within this class, whether they have taken place the past, or may take place in the future, or that we can imagine taking place.

Two-element set [*right*, *wrong*]

The two-element set $\{$*right*, *wrong*$\}$ is an abstract construction. It is a dyad, a pair of things, but the two things themselves are merely symbols. We assign absolutely no meaning at all to the two words *right*, *wrong*—they are no more than arbitrary, abstract tokens that we may associate with actions.

Later we will link these token words to what people commonly mean when they use these same words, but for the moment no such linkage exists.

A mapping

A mapping is a rule of assignment, that associates each element in some set of things (called the *domain*), with a particular element in some other set of things (called the *codomain*). A proper mapping must specify the assignment for every element of the domain, and it must assign it to a single, unique element within the codomain.

In the case of an abstract morality, the domain is the *set of actions*, and the codomain is our two-element set $\{$*right*, *wrong*$\}$. So for each action in our set of actions, we want to assign either the token word *right*, or the token word *wrong*, according to some scheme. We call the token assigned to each action the *value* of the action according to that scheme. The scheme must explicitly specify the value of *every* action in our domain of actions—none of them can be omitted or left unaccounted for. And it must specify the value of each action *unambigously*—it can't be both *right* and *wrong* at the same time, or a little bit of both. It has to be definitively one, or the other.

If the assignment scheme satisfies both these requirements it is a proper mapping, and therefore by our definition it is a valid abstract morality.

So the basic idea of an abstract morality is that for every action in our entire domain of actions, no matter what, from double-parking to genocide, the morality assigns the action a value—either the value *right*, or the value *wrong*.

2.1 Linguistic clarity

Our dyad of token words may look the same as the common English words "right" and "wrong"; they may be spelled and pronounced the same way, but they are not the same. The words "right" and "wrong" are part of natural language, fully loaded with a multiplicity of shifting and sometimes ill-defined conventional meanings. The token words *right* and *wrong* on the other hand have no meaning, being merely elements in the codomain of an abstract mapping.

Clearly, there is the potential for confusion here. Throughout this development we will be using both dyads, and it is important that we are always clear about which we are using, whether the natural language terms, or the abstract terms.

To distinguish the two pairs of words clearly, we will italicize the abstract tokens *right* and *wrong*, as we have been doing thus far. We will write the natural language terms "right" and "wrong" unitalicized, and frequently within quotation marks, as in this sentence.

There is a similar issue with regard to the word "morality." Throughout this work we will be using the term "a morality" in a way that is quite different from the ordinary everyday meaning of the word "morality."

We have defined a new concept, and coined the term "abstract morality" to refer to it. So we own this new term, and no confusion can arise when we use it. But as this is central to our development we will use this term frequently, and for this reason will usually shorten it to just "morality."

But now we must be alert to the possibility of another linguistic ambiguity. The word "morality" is also part of real-world vocabulary, with its own multiplicity of complex and emotive meanings, and we will also have cause to refer to these real-world notions of morality.

It will usually be clear from context what we mean, but where there is any possibility of confusion we will explictly say either "abstract morality," or "real-world morality," to specify our frame of reference.

2.2 Analogy: a children's game

To make these abstract ideas a bit more real, let's consider a somewhat analogous situation. Imagine a couple of children who have invented a game. Let's say a couple of little girls, let's say Charlotte and Emily, let's say about six or seven years old. And they've made up a couple of nonsense words, and the nonsense words might be something like *juba*, and *miki*, and they apply one or the other of these two words to everything in their world. So cornflakes might be *juba*, oatmeal might be *miki*. Trees are *juba*, flowers are *miki*. Mummy is *miki*; Daddy is *juba*.

They are little girls, and they are discovering the power and the possibility of language. They are discovering that things have not just names, but also qualities, attributes, categorizations, taxonomies. And they have made the joyous discovery that they can create their own attributes if they wish, and partition their world accordingly.

Soon enough they will tire of this and move on to other creative play, but on this particular afternoon as we look in on them, this game absorbs their attention completely. And during its momentary tenure on their imagination, we can observe and think about what they are doing.

And so Charlotte and Emily map their world. They run around in their gingham dresses, and for anything they come across, according to some ineffable, inexplicable, childish logic of their own, some secret intuition they have, known to them but unknown to us, they can say whether it's *juba*, or *miki*.

The creation of our abstract moralities is in much the same spirit as Charlotte and Emily's game. Just as we have a domain, they have a domain. Our domain is the set of actions; theirs is the set of things they know. Our codomain is the two-element set $\{$*right*, *wrong*$\}$; theirs is the two-element set $\{$*juba*, *miki*$\}$. And both codomains are equally abstract, each consisting of a dyad of undefined tokens. And just like us they create a mapping. They look at their domain of interest, and to each element in that domain they assign a specific value, or token.

And from our point of view theirs is indeed a proper mapping, because apparently, it's defined on the entire domain of interest, and it's unambiguous—nothing is neither *juba* nor *miki*, and nothing is both at the same time.

If we want to know what value something has we can ask them, "What is pineapple?" and they can always tell us. If this has not already been mapped Charlotte and Emily will put their beribboned heads together and quickly agree upon the answer. These are little girls remember; their play is inherently cooperative. Perhaps we can trip them up, offering them something outside their domain of knowledge: "What is serendipity?" But no dice; they demand immediately to know what this new thing is, and armed with this knowledge they can tell us in a heartbeat, why of course, serendipity is *miki*.

Or perhaps they don't even need to know what it is—perhaps unbeknownst to us they are mapping things based on their phonetic properties, and serendipity, laden with sibilants and dentals, is an easy call.

(This is not so implausible. Though their made-up words are simple, two-syllable constructions, our creative young protagonists already know to invent words with wide phonetic separation. Clearly, these linguistically attuned girls are destined to become literary women).

2.3 Outcomes and moralities

As Charlotte and Emily create their mapping, possibly they may endow their nonsense words with some kind of meaning, or possibly not.

Perhaps if we observe we may soon twig what they are doing. Perhaps they are using some simple rule—they are just alternating say: one thing is *juba*, then the next is *miki*, then back to *juba* again. But not likely—this is not nearly interesting enough to hold these girls' attention. Or perhaps their mapping is completely arbitrary, the word assignments somehow chosen randomly, without any pattern or meaning. Again, not likely.

More likely then either of these is that they are doing something more creative, building their mapping with a more subtle mix of meaning and non-meaning. Perhaps as we observe we might think we have unravelled the mystery; it might seem that inorganic, hard, minerally things are *juba*—rock is *juba*, glass is *juba*, but softer, organic, femininy things tend to be *miki*, so we might speculate that *juba* is a bit like yang, and *miki* a bit like yin.

But there are inexplicable exceptions to this that make no sense at all, that are wholly impenetrable to our clumsy adult minds, and at the end of the day we are left scratching our heads. At the end of the day it is a game they are playing. The rules are theirs, to make and break at will, subject to no logic but their own private one.

But whatever meaning *juba* and *miki* may have, this meaning does not extend beyond Charlotte and Emily, and is in no way inherent in the two words themselves. This is clear because we could introduce the same game to two other little girls, and let them play independently. And almost certainly they will create a completely different mapping, with different meaning or non-meaning implied for the same dyad of token words.

Or indeed we could have a computer play this game, and not involve any more children at all. We could provide the computer with the entire universe of things known to children of that age, and program it to assign values to this domain randomly, and be confident that this will generate a mapping with no discernable meaning at all.

So it's clear that there are many different possible mappings, and a different one can be expected from each pair of girls, or from each running of our random computer program. We might call the result of playing the game an *outcome*, and so we can say there are many possible outcomes.

And so it is with our own game. Just as there are many possible outcomes to the juba-miki game, there are many possible abstract moralities we can define on our domain of actions. And no one of them is any better or worse than any other.

2.4 Real-world morality: definitions

Thus far we have been speaking in terms of our abstract framework, and here we can speak freely, because all our terms are fully defined. But in the next section we will turn our attention for the first time to the real world, and what we will call "real-world morality."

In the real world of people and society, the notion of morality is tremendously complex and confusing. The very term itself refers to multiple distinct concepts, and these concepts of morality are themselves subject to conflicting ideas about what, fundamentally, they really mean, and how best to describe and define them. And all this is for good reason: the notion of morality is an exceedingly elusive and intractable intellectual concept.

However, this confusion at large need not concern us—our only obligation is to define our own terms clearly.

To make reference to the real world we will define three terms. First, we'll define what we will call an "individual moral sensibility." Then we'll extend this idea to groups of individuals and define a similar concept called a "group moral sensibility." Finally, with these two definitions in hand we will define a "real-world morality," applying to both individuals and groups. All these are our own definitions, not necessarily conforming to any other conventional meanings or connotations these terms might carry.

2.4.1 Individual moral sensibility

We define an **individual moral sensibility** as an innate sense of aversion by an individual to certain kinds of behaviour, or to the commission of certain acts. This sense of aversion may be experienced with regard to behaviour and acts by the individual itself, or it may be experienced with regard to behaviour and acts by others.

So as we have defined it, moral sensibility is an internal subjective experience, just like any other kind of sensibility, and ultimately unknown to all but the individual itself.

But we note that this sensibility typically has observable external manifestations. The sensibility may be manifested by the avoidance by the individual of acts or behaviour that offend the sensibility. Or it may be manifested by explicit verbal expression of the sensibility; in particular, by the application of conventional natural language terms such as "right" and "wrong" to acts or forms of behaviour that do or do not provoke the sense of aversion. The sensibility may also be manifested as a compulsion to direct censure against others who commit offending acts or behaviour.

So this is a purely descriptive definition—we are simply putting a name to something we can plainly observe: a particular aspect of human, or other, emotional functioning.

2.4.2 Group moral sensibility

As a subjective experience, the notion of an individual moral sensibility technically applies only to individuals. But we can readily define a similar notion for associated groups of individuals such as societies, cultures, religions and nations.

We note that within associated groups of individuals there is typically a high degree of commonality or overlap among the moral sensibilities of the individual group members, and the scope of this commonality is generally known throughout the group. This naturally gives rise to a cultural expression of sensibility, where a particular formulation of individual sensibilities, based on the dominant majority sensibilities within the group, is understood to be representative of the group as a whole. We'll call this representative sensibility a **group moral sensibility**.

And as with individuals, this community sensibility also has clearly observable manifestations. These may include:

- Relative rarity of acts or behaviour within the community that offend the sensibility;
- Explicit verbal expression of the sensibility, in particular:
- Application of conventional natural language terms such as "right" and "wrong" to acts or forms of behaviour that do or do not offend the sensiblity;
- Strong forces of dissuasion against offending acts or behaviour. These include the imposition of peer pressure towards conformity to the sensibility, expressions of censure against violations, and for severe violations, imposition of explicit sanctions or punishment;
- Explicit codifications of the sensibility such as religious behavioural doctrine, or legal systems of permissions and prohibitions.

Again this is a purely descriptive definition. We are doing no more than putting a name to an observable aspect of cultural expression and functioning.

For convenience we will use the term "moral sensibility" to refer to either an individual or a group moral sensibility.

2.4.3 Real-world morality

We are now ready for the main event: our definition of a real-world morality, which is what we've been after all along.

Informally, we want to define this as *a complete characterization of an underlying moral sensibility, in terms of all the acts that offend the sensibility, together with all the manifestations of the sensibility with regard to those acts*.

Formally, we define the **real-world morality** associated with a moral sensibility as a data set, consisting of the set of actions, and for each element in the set of actions, the set of all manifestations of the sensibility regarding that action.[1]

[1] What we are really doing here is defining a real-world morality as a relation on the Cartesian product of the set of actions and the set of manifestations. But to provide proper definitions and explanations of these terms (relation, Cartesian product), and to justify the notion of a "set of manifestations," would take us too far afield and is out of scope for the intended audience. Instead we are using the term "data set" as an informal way of referring to the more rigourous concept of a mathematical relation, in the hope that this will be sufficiently clear and acceptable to the reader.

On this topic, we note that these various terms and ideas we are borrowing from the world of mathematics (set, mapping, domain, codomain, value) are technically not applicable in our present sphere of inquiry, since none of the things we are talking about is defined in true set-theoretic terms. Rather we are using these ideas as metaphors to help us maintain clarity and precision of thought. At the end of the day we are in the messy world of actions and consequences, not the pure one of mathematics.

So in effect we are defining a morality as *the complete cataloguing of the manifestations of a moral sensibility*. The cataloguing consists of the totality of acts or forms of behaviour about which the sensibility has anything to say, together with what it has to say about them. Essentially a morality is the constitution or makeup of a moral sensibility, expressed in terms of its observable manifestations.

Once more we note that this is a purely descriptive definition, stated wholly in terms of observable phenomena.

2.5 Duality between abstract and real-world moralities

Though our framework of abstract moralities is an artificial construct, divorced from and making no reference to the real world, clearly it strongly echoes what we see in the real world.

First there is the domain of definition we have chosen: the set of actions. And this is because in the real world it is actions that carry moral value. Since we are inventing something we are free to invent whatever we want, and for our domain of definition we could have chosen the set of vegetables, or the set of musical instruments, or whatever we felt like. But those domains are not of interest to us—we do not commonly think of potatoes or carrots or violins as being "right" or "wrong." It is actions that we apply our moral judgements to.

Next there is our choice of codomain: the dyad $\{right, wrong\}$. We could have chosen a codomain of any number of elements, not just two, and for those elements we could have used any symbolic representation we wished. But we didn't—we chose two token words that are very similar, indeed identical but for italicization, to the two natural language terms that are most commonly used to express moral value.

And finally, there is our use of the concept of a mapping to describe our abstract moralities, also echoing how moral systems are formulated in real life. Though we don't commonly think of real-world moralities in these terms, this is one way of thinking of them: an assignment of words expressive of moral judgement—such as "right" and "wrong"—to the universe of things we consider to be subject to such judgement.

Thus our framework of abstract moralities consists of some of the structural elements that we observe in "real-world morality," lifted out of the real world, and recast in the form of an idealized abstraction.

Furthermore, there is a duality between our abstract framework of abstract moralities, and the real world of real-world moralities. For any abstract morality we define, there is a real-world morality that may exist, or that we can imagine. And for any real-world morality, we can define a corresponding abstract morality. Thus every abstract morality has a real-world counterpart, or doppelganger, and *vice versa*.

This duality has limitations—many of the abstract moralities we can define do not, or could not, exist in the real world. But we can still imagine what they might look like if they did exist. Conversely many real-world moralities do not have true abstract morality counterparts, because they do not necessarily satisfy the technical requirements for a mapping—they may not include an unambiguous set of manifestations for every action. A real-world morality may be silent regarding some actions, and for others it may show manifestations of both "right" and "wrong" at the same time. Nevertheless we can define an abstract morality based on the unambiguous portion of the real-world morality.

We specify the counterparts in each direction in the obvious way. Given a particular abstract morality, we imagine a real-world morality where all actions that map to *right* under the abstract morality show the classic manifestations of "right," and all actions mapping to *wrong* show the manifestations of "wrong."

Conversely, given a real-world morality we define an abstract morality where all actions showing manifestations of "right" map to *right*, and all actions showing manifestations of "wrong" map to *wrong*. But we must also deal with the possible undefinedness or multi-definedness of a real-world morality. We do this by "completing" the morality as follows: those actions where the real-world morality is silent, we arbitrarily map to *right*; those actions where the real-world morality is ambiguous, we arbitrarily map to *wrong*.

This generates a true abstract morality that we can use as a working approximation to any real-world morality. Not ideal but the best we can do, since in our abstract world all moralities are true mappings, but in the real world they are not.

2.6 Examples

With the notion of an abstract morality and a real-world morality in place, let's now look at some examples.

2.6.1 The Amorality

The first abstract morality we'll look at is the one that maps everything to *right*. Every single action in our entire set of actions has the value *right*.

This is a strange kind of morality to define, but it's completely valid. It's valid because it's a proper mapping: it's defined for every point in our domain of actions, and for every one it's defined unambiguously. Note that this doesn't mean that *wrong* no longer exists—it still exists as an abstraction, it's still present in the codomain of our function along with *right*. It's just that *wrong* doesn't have anything mapped to it. The set of things in the domain that map to *wrong*, is the empty set.

We'll call this the Amorality. We call it this because under its real-world counterpart there are no moral prohibitions at all; there is no censure or sanction against anything. There is a complete absence of manifestations of "wrong" for all actions, no matter how much they might violate our conventional sense of outrage or impropriety.

Though we can imagine it, a true real-world Amorality is implausible. An individual or association can come close to Amorality, but not even the most criminally insane psychopath, or pathologically self-interested corporation is fully Amoral, since both object to some actions: namely the subset of actions that impinge against themselves negatively. Indeed, both react immediately and aggressively to any such action. A perfectly Amoral individual or organization must not even have any sense of impropriety for actions taken against itself.

2.6.2 The Puritan morality

The next morality we'll define is the one where everything maps to *wrong*. With apology to Puritanism for this false characterization, we'll call this the Puritan morality. We call it this because under its real-world counterpart all actions are prohibited. Just as under the Amorality nothing is considered "wrong," under the Puritan morality everything is.

A true real-world Puritan morality is even more implausible than a real-world Amorality, and would appear to be completely self-paralyzing. If we were somehow to conjure such a society into existence and set it in motion, then all actions, even quite benign ones, are subject to moral censure. So the moment anyone helps an old lady carry her groceries for example, well that would be considered improper on moral grounds, so the person would perhaps be subject to some kind of reprimand. But then the reprimanding person or persons would themselves be subject to reprimand, and so on, and within a matter of hours the entire society would be deadlocked. So it's hard to imagine anything actually happening at all in such a society.

We can note several things from these two examples. First they clearly illustrate the limitation of the abstract/real duality mentioned earlier; though these two moralities are valid as abstractions, their real-world counterparts do not actually exist. And this is true more generally; a valid abstract morality may be quite implausible in real terms.

Along the same lines we note that abstract moralities need not make any sense in real-world terms. Real-world moralities generally do make some sort of sense; they have a recognizable internal consistency, reflecting the patterns of individual moral sensibility from which they are ultimately derived. But an abstract morality need have no consistency or coherence at all. As long as it is a valid mapping we can define it according to whatever caprice we wish.

Lastly, note that the Amorality and the Puritan moralities are exact complements of one another—whatever maps to *right* in one maps to *wrong* in the other, and *vice versa*. And more generally, every abstract morality has a complementary morality, where the mappings are completely reversed. So whenever we introduce or define a new morality, we also get its complement for free.

2.6.3 Morality-J

With the largely hypothetical Amorality and Puritan morality behind us, let's now try to define a morality that might have some real-world viability. We'll start with perhaps the most fundamental principle of morality we can imagine: that causing harm to others is a bad thing. So on this basis we'll define Morality-J as follows: for every action, if any entity, anywhere, is harmed as a consequence of the action then we map it to *wrong*. Otherwise we map it to *right*.

We need not exert ourselves to imagine the real-world counterpart to this because in fact it exists: Morality-J is an almost exact dual to Ahimsa, the fundamental non-violence principle of the ancient Indian religion of Jainism. Ahimsa prescribes a strict practice of causing no harm or injury to any living thing, either directly or indirectly, whether by thought, word, or deed.

Ahimsa is perhaps the purest and most noble expression of morality of all, standing in stark constrast to our Western traditions of unlimited harm towards whatever living thing we please. By virtue of this principle Jainism is of all world religions perhaps the one most deserving of respect, and the world would be a vastly better place were something like it to prevail throughout all human society. Regrettably, it does not.

2.6.4 Morality-N

Though noble in intention, Morality-J and Jainism present some practical difficulties. Ahimsa is difficult to follow in practice—so much so that even within Jain society, it is only the Jain ascetics, the monks and nuns, who commit their lives to this fully. Also there are situations in life where there is no course of action, including taking no action at all, that does not cause some harm somewhere.

These observations are innocent enough. But we might allow a further observation of Morality-J, more sinister, to take root in our mind. We might note, correctly, that even though an action may cause harm in one place, there may be benefits to be gained in another. And perhaps, a demon might whisper in our ear, we should take this into account. That sounds reasonable enough, does it not? Especially if we happen to be both the agent of an action, and one of its beneficiaries.

And so we might define an abstract morality that incorporates both the element of harm to victims, and the element of benefits to perpetrators. There are many ways one could define a morality on this basis, but we will present a very simple example, defining Morality-N as follows: for every action, if the number of beneficiaries of the action exceeds the number victims we map it to *right*. Otherwise we map it to *wrong*.

Though very simplistic, this morality bears some resemblance to reality: there is an implicit premise in most real-world moralities that numbers make a difference—for example we tend to regard serial murder as more reprehensible than a one-off murder, and mass murder as more reprehensible still. And more generally we tend to regard crimes with larger numbers of victims as more reprehensible, and crimes with larger numbers of beneficiaries as less so.

Morality-N is fully valid as an abstract morality, but it has real-world consequences we may not care for. For example, individual rape maps to *wrong*, but now gang rape maps to *right*. This means that under its real-world doppelganger individual rape would be considered "wrong," but gang rape would not—on the contrary, this would be a morally acceptable action.

And indeed anything we conventionally consider to be a crime can be turned into something morally acceptable, simply by bringing in additional beneficiaries as participants, until there are more beneficiaries than victims.

Morality-N may therefore seem far-fetched in real-world terms. But it is not. On the contrary, as we will see later there are real-world moralities that incorporate precisely this numerical principle. And they impose on their victims harm vastly worse than gang rape, in order to benefit a numerical superiority of beneficiaries.

2.6.5 Group moralities

Thus far we've defined some abstract moralities and imagined their real-world counterparts; we'll now look at some examples of real-world moralities. As we'll see, essentially any group or individual represents a real-world morality.

Almost any group of associated individual exhibits manifestations of an underlying moral sensibility, and thus represents a real-world morality. In the case of religious groupings for example, religious doctrine typically includes a set of moral precepts to be followed by its adherents, and for large-scale or highly formalized religions this is expressed in the form of a written document or set of documents. By definition, such documentation constitutes a real-world morality.

In the case of Christianity for example, the defining documentation consists of the Bible, containing within its pages the various permissions and prohibitions of biblical doctrine. And this constitutes a real-world morality, because for any action it presents a set of manifestations—the manifestations in this case consisting precisely of whatever the Bible has to say about that action. Thus one could extract from its pages a complete cataloguing of every action for which it has anything to say, together with everything it has to say about it. And this, by definition, is a real-world morality.

As it stands this cataloguing does not define a true abstract morality because it's not a valid mapping; it's riddled with contradiction and ambiguity, containing numerous instances of both prohibition and permission for the same action, and many actions about which it has nothing definitive to say at all.

We can still extract an abstract morality from the biblical morality, and one way to do this would be to "complete" it in the arbitrary manner described previously. But it's more interesting to note that any particular interpretation of the Bible, fully resolving all ambiguities, also provides a valid morality.

Thus we could consult a particular theologian, and for any action of interest ask for his interpretation of the Bible. We can ask him to use the Bible as his guide, and if it's silent or ambiguous or self-contradictory or incomprehensible regarding the action, then it's up to him to sort it all out and resolve the ambiguities according to his own theological interpretation, and tell us the answer: is this action biblically "right," or "wrong"?

Perhaps for some particularly problematic actions even our theologian might be stumped, unable to make head or tail of what the Bible has to say despite his years of study. In this case we can complete his interpretation, just as we completed the literal, uninterpreted Bible. Alternatively we might take a more demanding approach: we might insist upon a definitive answer one way or another, and to ensure we get it no matter how thorny a question we ask, we might lock our theologian in a room and assure him he won't be coming out until he has an answer for us.

But either way, this now gives us a valid abstract morality: precisely the real-world Biblical morality, resolved into an unambiguous mapping with the help of our theologian. And of course each distinct theological interpretation of the Bible provides a different abstract morality.

These observations are true of religious doctrine in general: the doctrine by definition constitutes a real-world morality. This is typically not a true abstract morality; but any complete and non-ambiguous theological interpretation of the doctrine does qualify as a valid abstract morality.

And though we've chosen religion and religious doctrine as an example of a group morality, clearly the same observations apply to *any* group exhibiting a moral sensibility, and for any expression of the sensibility: the expression constitutes a real-world morality by definition. And we can extract from this a corresponding abstract morality either by arbitrarily completing the ambiguous portions, or by seeking out from within the group an unambiguous "interpretation" of the expression.

2.6.6 Personal individual morality

These observations of group moralities apply similarly to individual moralities—indeed, we can think of an individual as nothing more than a group of one. Thus each individual, walking around with its own personal bag of sensibilities, is in effect a walking real-world morality.

And we can characterize the sensibility, and the associated morality, by quizzing the individual about any action of interest. We can ask, how do you feel about this action, say a woman's right to choice—are you personally OK with this, or no? And generally speaking anyone can eyeball this or any other action, and say how they feel about it.

As with groups, the personal individual morality is likely not to be perfectly unambiguous. We can expect the ambiguities for individuals to be less than for groups, but even individuals are likely to be conflicted about those actions lying in the anguished no-man's-land between clearly "right" and clearly "wrong." But as for groups, we can extract the corresponding abstract morality either by completing the personal morality, or by pressing the individual to give us his best-effort answer despite the conflicts.

And since we can do this with anyone claiming any sort of moral sensibility, we see that there are billions of personal individual moralities walking around. And perhaps like snowflakes, with no two exactly the same.

2.7 What we've accomplished

So what have we really accomplished here? On the face of it, not very much. We've defined the notion of an abstract morality, a completely made-up thing, and given some examples.

But in fact we've accomplished something quite important: we've chosen *a way of looking at things*. And for many difficult problems this is key—choosing the right way of framing the problem, so that the key issues can be seen with clarity.

We've chosen to look at real-world morality in purely empirical terms: we see that people have a strong moral sensibility, and we see clear manifestations of this among both groups and individuals. And we've chosen a particular way of characterizing this phenomenon: we think of it as a relationship, or mapping, between the things that provoke the sensibility (actions), and the expression of the sensibility (manifestations). We have then created an idealized representation of this characterization, in the form of our framework of abstract moralities.

This idealized framework allows us to maintain clarity of mind. The real world is confused and contentious. And any discussion of morality framed entirely in real-world terms is fully exposed to the turbulence of real-world passion and prejudice, to the many conflicting ideas and beliefs about the meaning of morality.

But within our abstract world there is none of that. Here there is only the serenity of logic. Here we can think and analyse without having to deal with any of the contradictions and conflict of the real world. Within our abstract framework we are all on the same page, regardless of our real-world beliefs—the Amorality maps late-term abortion to *right*, and no matter how you might feel about this action in the real world there is no arguing about it, because *that's the way we defined it*.

Within our abstract world there is no arguing at all—there is only understanding or not understanding.