

Ethical Theory and Technology

Jonathan Y. Tsou and Kate Padgett Walsh

Are automated driving cars good for society? To what extent should companies be able to profit from patents for developing innovative technological inventions? Should the use of cognitive enhancing drugs be allowed? Normative ethical theories provide different philosophical reasons for supporting conclusions on these and other ethical issues concerning technology. This chapter provides an overview of influential normative ethical theories and how they can be applied to particular cases in the ethics of technology.

Normative ethics addresses questions concerning what is morally good or right. Normative ethical theories (e.g., utilitarianism, Kantian ethics, contractarianism) specify different *criteria for moral rightness and wrongness* and how to live a morally good life by appeal to alternative philosophical principles. *Meta-ethics* addresses fundamental philosophical questions concerning the existence and nature of moral norms and judgments (e.g., are moral standards universal or culturally relative? Do moral facts exist? Are there moral truths?). By contrast, normative ethical theories assume that some moral standards (e.g., happiness, equity, freedom) are universal and articulate criteria for moral goodness based on these assumptions. *Applied ethics* is the application of normative ethical theories to specific cases (e.g., abortion, capital punishment).

Normative ethical theories are typically grouped into three categories: consequentialism, deontology, and virtue theory:

Consequentialism	Deontology	Virtue Theory
Utilitarianism	Kantian Ethics	Virtue Ethics
	Pluralism	
	Social Contract Theory	

Consequentialists maintain that moral goodness is determined exclusively by the consequences of action (or rules) and that there is nothing *intrinsically* good or bad in action itself. By contrast, deontologists claim that we are morally bound by certain moral duties and that some actions are inherently right or wrong. For deontologists, the morality of action is determined by whether

action conforms to specific moral duties (or rules). Rather than focusing on the consequences of action or following moral rules, virtue theorists assert that moral goodness is determined by whether our actions stem from virtuous (i.e., excellent) character traits.¹

Utilitarianism

Utilitarianism holds that an action (or rule) is morally right if it leads to the best consequences (greatest good) for the greatest number of people. Accordingly, the moral course of action is to promote the general good by maximizing utility, where ‘utility’ is the net good created by an action (or rule). In principle, utility can be defined in terms of various human goods (e.g., happiness, health, freedom). The classic utilitarians, Jeremy Bentham ([1789] 1907) and John Stuart Mill ([1861] 1998), defined utility in terms of happiness or pleasure (Driver, 2014). This is encapsulated by Bentham’s ‘fundamental axiom’ of utilitarianism: “the greatest happiness of the greatest number . . . is the measure of right and wrong” (cited in Crimmins 2021).² Compared to *ethical egoism*, which holds that we should act in our own self-interest (Shaver, 2021), utilitarianism holds that we ought to act in the best interests of everyone.³ Utilitarianism is a consequentialist theory because moral rightness is determined *exclusively by the consequences of action*; there is nothing *inherently* good or bad about *action itself*.

Act Utilitarianism (AU) and Rule Utilitarianism (RU) are two forms of utilitarianism that focus on the consequences of actions and rules, respectively:

Act Utilitarianism (AU): actions that maximize utility are morally good.

Rule Utilitarianism (RU): rules that maximize utility are morally good.

In the philosophical literature, AU and RU are commonly presented as *competing ethical theories* that provide conflicting assessments on the *morality of action*. For example, J. J. C. Smart (1956) presents AU as the theory that moral action is an action that maximizes utility in a specific circumstance, whereas RU holds that moral action is an action that corresponds to a rule (e.g., “keep promises”) that tends to maximize utility. Based on this distinction, Smart argues that it is irrational to endorse RU because it implies that agents should perform an action (e.g., keeping a promise to a dying acquaintance to donate their money to an unsavory institution) even if another course of action (e.g., donating to a charity) would obviously be better for maximizing utility.

While much of the philosophical literature assumes that AU and RU are competing theories that prescribe different normative standards *for evaluating the morality of action*, AU and RU can also be understood as *complementary theories* that have different contexts of application: AU is a theory that assesses the morality of action in specific circumstances,

¹ Since the revival of virtue ethics in the 20th century, it has been commonplace to distinguish ethical theory and moral theory in terms of their scope of application (e.g., see Williams, 1985). In this chapter, we follow the distinction that ethical theories are concerned (broadly) with the goodness or badness of the agent's character or dispositions, whereas moral theories are concerned (more narrowly) with the rightness and wrongness of actions.

² Bentham and Mill are ‘hedonistic Utilitarians’ who assume that pleasure is the only intrinsic good and pain is the only intrinsic bad (Sinnot-Armstrong, 2021). While Bentham maintains that pleasure and pain only differ *quantitatively* (e.g., in terms of intensity and duration), Mill argues that there are *qualitative* differences between lower-level pleasures (e.g., sensual pleasures, drunkenness) and higher-level pleasures (e.g., intellectual achievements, spiritual gratification) (Driver, 2014; cf. Crimmins, 2021).

³ Mill ([1861] 1998) held that this axiom can include non-human animals, since they are sentient beings who experience pleasure and pain. Peter Singer ([1975] 2009) famously argues for animal rights on the Utilitarian ground that non-human animals experience pleasure and pain (e.g., suffering).

whereas RU is a theory that assesses the morality of general rules (e.g., proposed laws or policies). As such, AU and RU need not be regarded as conflicting theories, but complementary *applications of utilitarian theory*.⁴ Applied ethics is typically less interested in evaluating the morality of specific actions and more interested in evaluating the morality of general rules (e.g., policies, laws) regarding a *class of ethical cases* (e.g., abortion, euthanasia, capital punishment). Whereas AU can assess the morality of action in a specific context (viz., prefer the action that maximizes utility), it provides no specific guidance for *evaluating the morality of general rules*. On this issue, RU maintains that rules (e.g., laws or policies) that maximize utility are morally good.

In ethical cases involving technology, RU can be regarded as an application of utilitarian theory that assesses the morality of policies regarding technology. For instance: Should society tolerate privacy technologies that can virtually secure anonymity online? Should society provide financial incentives for producing electric vehicles? There *are* questions in technology ethics about the morality of a specific action (e.g., was Julian Assange morally justified in publishing confidential US military information on WikiLeaks?). But many questions concern the morality of more general policies (or rules) related to technology (e.g., should democratic societies tolerate media sources that publish classified or confidential information for the sake of public transparency?). RU can be regarded as an application of utilitarianism that assesses the morality of rules in terms of the positive and negative (actual or possible) consequences of adopting such rules.

A strength of utilitarianism it provides a concrete method for determining the moral course of action. For a given case, it directs us to operationally define utility (e.g., happiness) and weigh the positive consequences against negative consequences. For cases of policy (e.g., capital punishment), this requires comparing the possible gains of adopting a policy (e.g., providing a deterrent for heinous crimes, vindication for families of victims, lowering costs of incarceration) against possible losses (e.g., the wrongful killing of innocent suspects). In principle, utilitarianism thus offers a straightforward method for determining the moral course of action. In practice, applying utilitarian theory is more complex. There are alternative ways in which “utility” can be defined and utilitarianism does not provide explicit guidance on how to weigh conflicting positive and negative utilities. In this regard, there is a tradeoff in how specifically utility is defined and how readily applicable utilitarian theory is. If one assumes a simple definition of utility (e.g., happiness), utilitarianism is more easily applied over a broad range of cases, but it will gloss over nuances regarding utilities (e.g., negative utilities that could assigned greater value). If one assumes a more sophisticated definition of utility (e.g., a view that places greater weight on certain types of pleasure or pain), then it can better account for nuances, but will be more difficult to apply mechanically over a broad range of cases.

One weakness of utilitarianism is its insensitivity to individual rights. For example, if we adopt a simple definition of utility, wherein each individual human life is assigned a utility value of +1, utilitarianism implies that sacrificing the life of two people (-2) to save the lives of eight people (+8) is not only morally permissible, but morally required (net utility of +6). On this view, it would be acceptable to murder two healthy individuals if their organs could be used to save the lives of eight patients who require organ donation. Similarly, it would be acceptable for a family stranded at sea to eat two family members to save the lives of eight. These somewhat extreme counterexamples are exemplified by ‘trolley case’ thought experiments, which analyze

⁴ This is arguably Mill’s view. For discussion of various interpretations of utilitarian theory, see Sinnott-Anderson (2021).

the ethics of sacrificing innocent lives in order to save a larger number (Wollard & Snyder, 2021). Another drawback of Utilitarianism is its neglect of *intentions* and character. Utilitarians regard an action (e.g., creating a vaccine for a deadly disease) as morally good even if the action was motivated entirely by bad reasons (e.g., the inventor acted entirely out of selfish and greedy motives). These considerations suggest there are other relevant factors (other than consequences) that matter for ethics: certain actions (e.g., violating the rights of others) are *inherently unethical*. While utilitarians hold that there is nothing besides possible or actual consequences that matter for ethics, deontological theories (e.g., Kantian ethics, social contract theory) aim to articulate the theoretical principles that make our actions inherently moral or non-moral.

Kantian Ethics

Immanuel Kant argues that *why* we act is essential to the morality of our actions. Regardless of the consequences of actions, what matters for ethics is our *motives* (or intentions). For Kant, a person who only helps others for self-interested reasons is not truly acting morally. For example, consider the different motives one can have for saving a stranger who is drowning. When we help strangers independent of our own selfish motives (e.g., receiving a reward), we are motivated, Kant says, by good will: the intention to do the right thing just because it is the right thing to do. Good will differs from other motives because it is not self-interested. To translate good will into action, we need a moral principle (or test) for determining the moral goodness and badness of action. Kant calls this principle the ‘Categorical Imperative’ (CI). CI holds that we have a duty to follow moral rules that could be applied *universally and effectively* by all people in the same sort of situation.

For Kant, the CI is the fundamental moral principle from which all other moral duties are derived. The CI is an ‘imperative’ because it is a command (or rule) that rational agents have a moral obligation or duty to follow. It is ‘categorical’ because it prescribes moral duties that apply in all situations without exception. Kant formulates the CI in various ways (Johnson & Cureton, 2021). Kant’s first (‘Universality’) formulation of the CI is perhaps the most important: “act only in accordance with that maxim through which you can . . . will that it become a universal law without contradiction” (Kant, [1785] 2012,4:421). This suggests that we should follow principles of action (‘maxims’) that could be *consistently and effectively applied as a normative moral rule* (‘universal law’) by all rational beings.⁵ Doing so requires that we first identify the general principles (e.g., tell the truth, do not harm others) that direct our action in particular circumstances. Second, consider whether these principles could serve as universal laws that all people follow. The test of the latter is whether we can conceive of a world where these principles are *consistently applied to everyone* (‘without contradiction’) *to achieve their intended purposes*.

Several moral duties (e.g., tell the truth, keep promises, do not harm others, help others) are clearly derivable from the CI. For these duties, we can conceive of worlds where, *if*

⁵ This formulation is complemented by Kant’s third (‘Autonomy’) formulation and the ‘Kingdom of the ends’ formulation of the CI (Johnson and Cureton, 2021), which both emphasize the universality of our moral duties and their status as potential *laws* that could legislate the morality of rational beings.

everybody followed these rules, they would achieve their intended purposes (e.g., supporting the integrity of communication, building social trust, promoting cooperation). In contrast, if an attempt to universalize a rule results in a contradiction, then acting on that rule is morally prohibited (Korsgaard, 1996). For example, consider a “promise breaking” principle: “keep promises except when it is inconvenient.” This principle cannot be moral duty because universalizing it results in a contradiction. In a world where everybody followed this principle, following this principle would not be practically effective (since ‘promises’ would have no credibility in such a world). For the principle of selective promise-breaking to be practically effective, there need to be others in the world who follow a promise keeping principle to give credibility to promises. What the CI prohibits is following rules that allow us to make an exception of ourselves.

Kant complements his universality formulation of the CI with the second (‘Humanity’) formulation of the CI: “[a]ct in such a way that you treat humanity, whether in your own person or in the person of another . . . as an end and never simply as a means” (Kant, [1785] 2012, 4: 429). This formulation suggests that we should never treat other humans *exclusively* as means to an end (i.e., as instruments for achieving our self-interested goals), but as ends in themselves (i.e., as individual persons who possess *inherent value*). This formulation prescribes that all people should be respected *because of their inherent value as persons*. We should not treat persons just as objects or tools (e.g., slaves, child laborers). Doing so is immoral because it violates a basic respect that should be afforded to all persons. According to the second formulation of the CI, coercion and lying are thus paradigms of immoral behavior because they: (1) violate the rights of others to consent, and (2) treat other persons (or their beliefs) as mere instruments (Korsgaard, 1996).

A strength of Kantian ethics is that it provides an objective and impartial foundation (i.e., consistency of application) for identifying moral duties. For Kant, the morality of action is determined by whether our motives for action are consistent with the CI, which prescribes that we follow moral principles that can be applied fairly and effectively to govern all rational beings. When we act by principles that could not be followed by everyone, we act immorally because we make an exception of ourselves and exploit rules that we expect others follow. In the context of technology ethics, a Kantian might argue that internet trolling (i.e., participating in internet conversations with the sole intention of upsetting others) is wrong because violates our duties to be honest and to treat others (as emphasized in Kant’s second formulation of the CI) with basic respect.

A weakness of Kantian ethics is its rigidity. Because the moral duties derived from the CI are *categorical*, it is *always* wrong to violate duties prescribed by the CI. Since lying is *always wrong*, it is unethical to tell a lie even if it leads to good consequences (e.g., lying to an assassin to save an innocent person’s life, telling a small lie to spare a friend’s feelings). In such cases, different Kantian moral duties (e.g., tell the truth, help others) can appear to conflict. Kantians suggest that such conflicts can be resolved by more precisely specifying the principles of our proposed actions so as to include morally relevant details (see Johnson & Cureton, 2021), but the difficulty remains. Relatedly, in its emphasis on universalization, Kantian ethics can sometimes

seem overly rationalistic, and thus less sensitive to the particular contexts of ethical cases. In applied ethics and technology ethics, we are often interested in evaluating complex principles that involve numerous conflicting moral duties. For example, questions concerning whether society should tolerate privacy technologies (e.g., encryption, TOR) that can virtually guarantee anonymity online can be reduced to the question of whether the duty to respect everyone's privacy or the duty of governments to protect its citizens (since privacy technologies can be exploited by criminals) is more important. In complex cases, attempts to universalize the relevant principle will not always indicate clearly whether such imagined worlds imply a contradiction.

Ross's Pluralism

An alternative deontological theory is William Ross's intuitionist approach to ethics, which maintains that moral goodness involves following a plurality of moral duties. Unlike Kant, who derives our duties from a fundamental moral principle (i.e., the CI), Ross assumes that our moral duties are fairly self-evident and recognizable ('intuited') by rational agents. Ross conceptualizes our moral intuitions as initial 'intellectual appearances': initial reactions or impressions that seem self-evident (Huemer, 2005). Ross contends that our intuitions reveal a plurality of 'prima facie' moral duties: moral duties that we have an obligation to follow *most of the time*.

Ross' pluralism suggests that we have a number of prima facie moral duties that we are ethically required to follow:

1. Fidelity: keep promises and tell the truth.
2. Reparation: correct wrongdoings that you have inflicted on others.
3. Gratitude: return services to those who help you.
4. Non-injury (Non-maleficence): do not harm others.
5. Beneficence: be kind to others and to try to improve their well-being (e.g., health, wisdom, security).
6. Self-improvement: improve your own well-being.
7. Justice: distribute benefits and burdens equably and evenly.

In presenting these duties as prima facie obligations, Ross ([1930] 2002) suggests that our moral duties are not categorical duties as Kant assumes, but duties that we should follow most of the time.⁶ Moreover, in contrast to Kantian ethics, Ross's moral duties are not organized hierarchically and, in cases where duties conflict, *the context of a case* will determine which moral duty is most salient.

A strength of Ross' pluralism is that it provides a compelling list of moral duties that most people would agree are good moral rules to follow. In this regard, it is worth noting that Ross' duties are derivable from Kant's CI *and* they are principles that a Rule Utilitarian would endorse as good moral rules. Compared to Kantian ethics, a strength of Ross' theory is that it more directly addresses the particular circumstances surrounding ethical cases. For example, in

⁶ Ross ([1939] 2019) suggests that 5-7 could be grouped together as a general duty to maximize the overall good (i.e., to promote values that are intrinsically good).

cases where the moral stakes are not particularly high (e.g., a romantic partner asks if they look attractive in a certain outfit), Ross' pluralism would tolerate small lies (i.e., the duty of beneficence outweighs the duty of fidelity). In cases with higher moral stakes (e.g., an assassin asks you whether an innocent person is hiding in your house), pluralism suggest it is morally permissible to lie (i.e., the duty of non-injury outweighs the duty of fidelity).

A weakness of Ross' pluralism, compared to Kantian ethics, is that moral duties are not systematically derived. This raises the possibility that when applying pluralism to specific ethical cases, one is simply reaching ethical conclusions in an *ad hoc* manner (i.e., using pluralism to justify moral conclusions that we want to reach prior to analysis). Relatedly, it will often be difficult to determine the moral course of action in cases where duties conflict. For example, when hiring job candidates, there will be cases where the duty of justice (e.g., offer the job to the best candidate) and the duty of beneficence (e.g., offer the job to the candidate who will benefit the most) conflict. In such cases, Ross' theory offers no specific guidance on how to settle such conflicts besides looking at the context. Another weakness of pluralism is its reliance on moral intuitions. In recent years, empirical and experimental research into philosophical methodology has suggested that intuitions—including moral intuitions—are an unreliable source of philosophical evidence (Knobe & Nichols, 2017; Pust, 2021). For example, intuitions about prominent ethical thought experiments show systematic variation across cultures, which suggests that our moral intuitions are determined (or at least strongly influenced) by cultural factors (e.g., moral norms prevalent within a culture). This speaks against Ross' assumption that rational agents can reliably 'intuit' our moral duties and obligations.

Social Contract Theory

More recent deontological approaches draw on social contract theory. Historically, social contract theory was defended in the 17th and 18th centuries by Thomas Hobbes, John Locke, Jean-Jacques Rousseau, and Kant as a political theory. It was subsequently developed in the 20th century as a normative moral theory. Social contract theorists aim to provide a rational justification for normative ideals (e.g., justice, moral goodness) by appeal to the idea of a social contract: our political and moral obligations are ultimately justified by *an agreement among rational, adequately informed agents to follow certain contractual rules*. As a moral theory, this implies that our moral duties are not Kantian universal duties nor Rossian intuited prima facie duties. Instead they are *contractual duties* that rational agents would mutually agree or consent to. Two distinctive approaches in contemporary social contract theory are *contractarianism* (which follows a Hobbesian tradition) and *contractualism* (which follows a Kantian tradition). These approaches diverge on what grounds rational agreement regarding contractual duties. Whereas contractarians ground this agreement in Hobbes' idea of mutual self-interest, contractualists ground rational agreement in Kant's ideal of the equal moral status of persons (Ashford & Mulgan, 2018; Cudd & Eftekhari, 2021). Prominent contemporary examples include Rawls' contractualist political theory, Gauthier's contractarian moral theory, and Scanlon's contractualist moral theory.

John Rawls is the most important contemporary social contract theorist, and his theory of justice (Rawls, [1971] 1999) spurred a revival of social contract theory in the 20th century. Rawls' theory is driven by the idea that our political obligations should be governed by a respect for individual rights and a philosophy of fairness. On this view, all individuals possess certain rights in virtue of an implicit social contract with society. For Rawls, this social contract is not actual, but a hypothetical ideal contract that rational individuals would agree to. Rawls considers what kind of contract individuals would consent to in the 'original position,' where individuals are behind a 'veil of ignorance' and have no knowledge about their personal characteristics (e.g., age, race, social standing) or the society. Rawls argues that rational agents would agree to a contract that *protect everyone's rights* equally: rules that protect people's rights equally are just. From this perspective, Rawls ([1971] 1999) defends a 'greatest equal liberty principle': "Each person is to have an equal right to the most extensive total system of equal basic liberties compatible with a similar system of liberty for all" (p. 266). Assuming that freedom and happiness are fundamental basic liberties, individuals have a right to maximize their freedom and happiness, so long as they do not infringe on others' right to these basic liberties. When the interests of groups conflict, Rawls ([1971] 1999) proposes his famous 'difference principle,' in which inequalities in the distribution of goods are justified only if they benefit the least advantaged members of society (p. 266). Rawls' theory implies that justice is fairness insofar as the hypothetical social contract is a *fair contract* that rational individuals would agree to. The veil of ignorance assures that agreement is based on the interests of everyone, rather than self-interest.

In the context of contemporary ethical theory, David Gauthier and T. M. Scanlon defend influential social contract theories that offer alternative justifications for our moral obligations. Gauthier (1986) defends a contractarian ('morals by agreement') theory that justifies our moral duties in terms of self-interest. Like Hobbes, Gauthier assumes a naturalistic view of humans (i.e., as self-interested creatures), an instrumental account of rationality (i.e., as a practical ability to meet agent-relative goals), and a conventionalist account moral duties (i.e., as rules that promote mutual self-interest). Gauthier maintains that we should be moral because it is in everybody's mutual interest to do so. If everyone followed moral rules (e.g., do not lie, do not harm others) that promote social cooperation and mutual benefit, then everyone would be better off. In cases where self-interest (e.g., earning more money) conflicts with moral action (e.g., paying taxes), morality demands that we *constrain* our self-interested pursuits. The justification and normative force of these constraints stems from their rationality: rational self-interested individuals would agree to constrain their behavior and follow moral rules that are mutually beneficial, so long as others also follow these rules.⁷ Scanlon (1998) defends a contractualist theory that justifies our moral obligations in terms of the Kantian ideal (expressed in Kant's Humanity formulation of the CI) that all rational agents should be valued and respected *equally*. According to Scanlon (1998), an act is wrong if it would be prohibited by moral principles that *no rational agent could reasonably reject* (p. 153). An important feature of rational agents is

⁷ Gauthier explicates the (practical) rationality of moral agents in the language of rational choice theory. He presents rational moral agents, alternatively, as 'constrained maximizers' (Gauthier, 1986) or 'rational optimizers' (Gauthier, 2013).

their capacity to recognize and act upon justified reasons. Hence, treating other rational agents as equals requires acting towards them in accordance with principles that they (and we) could not reasonably reject. On this view, actions wrong if they are *unjustifiable* to other rational agents. Whereas Rawls' contractualist theory invokes self-interest behind a veil of ignorance to establish agreements that account for everyone's interests equally, Scanlon's theory maintains that the impetus to respect everyone's interests equally is driven by an intrinsic desire *to justify our actions to others*. Possessing this kind of desire (akin to Kant's 'good will') is part of what it is to be a rational moral agent (Ashford & Mulgan, 2018).

Social contract theories offer useful resources for evaluating complex ethical cases where the interests of different groups conflict. Such theories emphasize identifying and justifying moral (contractual duties) that are *fair* and *impartial*. For example, questions concerning the regulation of encryption and other privacy protection technologies (e.g., TOR, cryptocurrency) concern a conflict between the right of citizens to protect their privacy (e.g., against governmental surveillance) and the right of governments to regulate what technologies are widely available to its citizens (e.g., to protect public safety). One could give a Rawlsian argument that the rights of individuals should take priority in this case because they are the larger and more disadvantaged group. Alternatively, one could appeal to Gauthier's contrarian analysis to evaluate whether a policy (e.g., that the government should have some power to regulate privacy technologies) is fair in the sense of advancing everyone's mutual self-interest. Given the specific questions that social contract theories are formulated to address, one limitation they face is their applicability. Rawlsian analyses are most relevant to issues of social justice and the state within liberal and democratic societies, which is an important, but narrow aspect of ethics. Scanlon's contractualism is most relevant for questions concerning our moral obligations to one another.

Virtue Ethics

In contrast to consequentialist and deontological theories, virtue ethics focuses on understanding how we should live our lives and what sorts of persons we should be. Compared to the moral theories reviewed above, which evaluate the rightness or wrongness of action (see note 1), virtue theorists address broader ethical questions concerning how to live the good life. According to the Greek philosopher Aristotle and the Chinese philosopher Confucius, the ultimate goal of ethics is to cultivate a good character and live a flourishing life, rather than just to decide which actions are right and which are wrong. In the second half of the 20th century, virtue ethics was revived by philosophers who were disenchanted with the narrow focus on actions within consequentialism and deontology.

Virtue ethics starts with the idea that humans, like all living beings, have lives that can go well or badly. To flourish rather than languish, we need good nutrition and opportunities to be physically active, just like other animals. But because of our complex cognitive abilities and social natures, we also need to have a *strong personal character* to live well. And a strong character is made up of many different virtues (e.g., courage, generosity).

Aristotle (1984, books 2-4) discusses a number of moral virtues, which represent an optimal point ('mean') between two vicious extremes. Some of the virtues proposed by Aristotle

include the following:

Vice (deficiency)	Virtue (Mean)	Vice (excess)
Cowardice	Courage/ bravery	Rashness/ recklessness
Self-indulgence	Temperance/ moderation	Abstinence/ insensibility
Stinginess	Generosity	Extravagance/ prodigality
Undue humility	Pride	Vanity
Indifference	Moderate temper	Irascible/ easily angered
False modesty	Truthfulness	Boastfulness
Bashful/ Ashamed	Modesty	Shameless
Quarrelsome	Friendliness	Obsequious/ fawning
Rudeness/ boorishness	Wittiness/ charm	Foolishness/ buffoonery

Aristotle’s virtues apply to specific domains (‘spheres’) of feeling or action. For example, courage applies to the domain of fear and confidence; generosity to the domain of spending money; and wittiness to the domain of social conduct. In this theory, virtues represent an *optimal* (‘mean’) amount of some character trait that falls between two extremes. For example, someone who has too little courage is a coward, but someone who is courageous to the point of lacking caution is reckless. The rightness or wrongness of actions is, then, determined by whether a person acts in accordance with virtuous or vicious characteristics. When someone drinks excessive amounts of alcohol to get drunk, they act immorally because they act upon the vice of self-indulgence. When someone interacts rudely with others on social media, they act immorally because they act upon vices such as vanity, shamelessness, and excessive anger.

A strength of virtue ethics is that it offers concrete guidance on how to acquire strengths of character. If we can identify moral virtues that contribute to a flourishing life, then we can attempt to cultivate those virtues ourselves. These virtues are to be developed and improved through experience, rather than by learning a rule and simply applying it. We learn by following virtuous role models and by training ourselves to act virtuously. To become a more generous person, for instance, you could commit to looking for opportunities to donate to charities. Over time, by creating a habit of acting generously, you would slowly begin to think more about helping others and thereby acquire the virtue of generosity. For cases of technology ethics, e.g., a virtue theorist might analyze the ethics of social media use in terms of virtuous characteristics (e.g., moderation, pride, friendliness).

Like other ethical theories (e.g., Kantian ethics, contractarianism), a weakness of virtue ethics is that it is sometimes not obviously applicable to complex ethical cases. When trying to decide how to act, virtue ethics recommends that we do whatever a virtuous person would do;

however, in complex cases (e.g., stealing food to say a starving family), different virtues may push us towards conflicting behavior and virtue ethics does not offer a specific procedure of weighing different virtues against one another. Another difficulty concerns the nature of the virtues themselves. Among contemporary proponents of virtue theory, there is disagreement on what the correct moral virtues are and how virtues should be conceptualized (Hursthouse & Pettigrove, 2018). Moreover, given cultural differences, there is a worry that different virtues identified by virtue theorists might reflect *prevalent, culturally-relative social norms*, rather than universal virtues (MacIntyre, 1985). For example, among Aristotle's proposed virtues, it seems peculiar to argue that someone acted *immorally* if they were too shy or not witty enough in a social setting. Hence, various moral virtues proposed by philosophers might simply codify and reinforce culturally-relative norms (cf. Nussbaum, 1993). In addition, some have challenged virtue ethicists' assumption that character traits are sufficiently stable across time and situations to translate into moral behavior. A growing body of "situationist" psychological research suggests that many of our character traits are specific to situations (Doris, 2020). For example, someone might be generous whenever they are relaxed, but stingy when stressed out. If character traits are not stable across time and situations, then the project of cultivating virtues is much more complex than it initially appears.

Normative Ethics and Technology

The normative ethical theories discussed in this chapter offer different philosophical reasons for thinking that a particular action (or rule) about technology is morally good or bad. Importantly, these theories will often converge upon the same ethical conclusions albeit for different reasons.⁸ For example, all theories reviewed in this chapter reach the conclusions that murdering innocent people for fun is morally bad and that helping people in need is morally good. Where they differ is in the principled reasons that justify such conclusions. In the context of ethical cases concerning technology, different ethical theories provide alternative theoretical resources for evaluating and justifying ethical claims about actions or policies related to technology.

The broad summary of canonical ethical theories in this chapter provides a reference point for reading subsequent chapters in this volume. Other chapters provide more detailed examinations of the normative theories discussed, an examination of alternative normative ethical theories not discussed, and applications of such ethical theories to various cases in technology ethics.

⁸ Based on this convergence, Derek Parfit (2011) defends a moral theory that combines rule consequentialism, contractualism, and Kantian ethics.

References

- Alexander, Larry and Michael Moore. (2020). "Deontological Ethics." *The Stanford Encyclopedia of Philosophy* (Winter 2020 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2020/entries/ethics-deontological/>
- Aristotle (1984). *Nicomachean Ethics*. Translated by W. D. Ross; revised by J. O. Urmson. In *The Complete Works of Aristotle*, The Revised Oxford Translation, vol. 2. Edited by Jonathan Barnes. Princeton: Princeton University Press.
- Ashford, Elizabeth and Tim Mulgan. 2018. "Contractualism." *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/sum2018/entries/contractualism/>
- Bentham, Jeremy, (1789) 1907. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Crimmins, James E. 2021. "Jeremy Bentham." *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2021/entries/bentham/>
- Cudd, Ann and Seena Eftekhari. 2021. "Contractarianism." *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2021/entries/contractarianism/>.
- Driver, Julia. 2014. "The History of Utilitarianism." *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2014/entries/utilitarianism-history/>
- Franssen, Maarten, Gert-Jan Lokhorst, and Ibo van de Poel. 2018. "Philosophy of Technology." *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition). Edited by Edward N. Zalta, URL = <https://plato.stanford.edu/archives/fall2018/entries/technology/>
- Gauthier, David. 1986. *Morals By Agreement*. Oxford: Oxford University Press.
- Gauthier, David. 2013. "Twenty Five On." *Ethics* 123(4): 601-624.
- Hooker, Brad. 2016. "Rule Consequentialism." *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule/>
- Hursthouse, Rosalind and Glen Pettigrove. 2018. "Virtue Ethics." *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2018/entries/ethics-virtue/>

- Johnson, Robert and Adam Cureton. 2021. "Kant's Moral Philosophy." *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/spr2021/entries/kant-moral/>
- Kant, Immanuel. (1785) 2012. *Groundwork of the Metaphysics of Morals*, 2nd ed. Translated and edited by Mary Gregor and Jens Timmerman. Cambridge: Cambridge University Press.
- Knobe, Joshua and Shaun Nichols. 2017. "Experimental Philosophy." *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2017/entries/experimental-philosophy/>
- Korsgaard, Christine. 1996. *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.
- Kraut, Richard. 2018. "Aristotle's Ethics." *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/sum2018/entries/aristotle-ethics/>
- Lyons, David. 1965. *Forms and Limits of Utilitarianism*. Oxford: Clarendon Press.
- MacIntyre, Alasdair. 1985. *After Virtue*. 2nd ed. London: Duckworth.
- Mason, Elinor. 2018. "Value Pluralism." *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/spr2018/entries/value-pluralism/>
- Mill, John Stuart (1861) 1998. *Utilitarianism*. Edited by Roger Crisp. New York: Oxford University Press.
- Nussbaum, Martha C. 1993. "Non-Relative Virtues: An Aristotelian Approach." In *The Quality of Life* (pp. 242-270). Edited by Martha C. Nussbaum and Amartya Sen. Oxford: Oxford University Press.
- Parfit, Derek. (2011). *On What Matters*, Oxford: Oxford University Press.
- Pust, Joel. 2019. "Intuition." *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/sum2019/entries/intuition/>
- Rawls, John. (1971) 1999. *A Theory of Justice*, 2nd ed. Cambridge, MA: Harvard University Press.
- Ross, W. David. (1930) 2002. *The Right and the Good*. Philip Stratton-Lake, ed. Oxford: Oxford University Press.
- Ross, W. David (1939) 2018. *Foundations of Ethics*. Oxford: Clarendon Press.
- Scanlon, T. M. 1998. *What We Owe to Each Other*, Cambridge, MA: Harvard University Press.

Shaver, Robert. 2021. "Egoism." *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2021/entries/egoism/>

Singer, Peter. (1975) 2009. *Animal Liberation*. New York: Harper.

Sinnott-Armstrong, Walter. 2021. "Consequentialism." *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/fall2021/entries/consequentialism/>

Smart, J. J. C. 1956. "Extreme and Restricted Utilitarianism." *Philosophical Quarterly* 6(25): 344-354.

Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.

Woollard, Fiona and Frances Howard-Snyder. 2021. "Doing vs. Allowing Harm." *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition). Edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/fall2021/entries/doing-allowing/>