

## **Reformulating Mill's Harm Principle**

**Ben Saunders, University of Stirling**

*Mill's harm principle is commonly supposed to rest on a distinction between self-regarding conduct, which is not liable to interference, and other-regarding conduct, which is. As critics have noted, this distinction is difficult to draw. Furthermore, some of Mill's own applications of the principle, such as his forbidding of slavery contracts, do not appear to fit with it. This article proposes that the self-regarding/other-regarding distinction is not in fact fundamental to Mill's harm principle; what he should have said is that intervention is permissible only to prevent non-consensual harm, regardless of where it falls. This explains both why some other-regarding conduct is immune to interventions and why some self-regarding conduct can be interfered with.*

Keywords: Consent; Harm Principle; John Stuart Mill; *On Liberty*; Paternalism; State Regulation.

The so-called 'harm principle', set out in J. S. Mill's *On Liberty* has, despite its ambiguities, proved an enduring and influential contribution to debates over the limits of legitimate state/social action. Mill illustrated the principle through discussion of its potential applications to the regulation of, amongst other things, discussion and alcohol, and more recent writers have extended a Millian approach to new problems, such as genetic engineering (Holtug 2001; Burley 2008). Whether or not we ultimately endorse such an approach, however, it is imperative that we base our assessment on a proper understanding of Mill's principle. It is commonly supposed,

first, that the harm principle rests upon a distinction between self- and other-regarding actions, even though Mill does not use the phrase ‘other-regarding’ himself (Wollheim 1973, p. 2; Riley 1991, pp. 15-6) and, second, that this distinction is problematic. This article rejects the first assumption, thereby showing that Mill, and contemporary defenders of the harm principle, need not be troubled by well-known problems with identifying a purely self-regarding sphere of action.

The traditional interpretation is not without textual basis. Mill’s opening statement of his harm principle distinguishes between harm to others and harm to oneself. “[T]he only purpose for which power can be rightfully exercised over any member of a civilized community, against his will, is to prevent harm to others. His own good, either physical or moral, is not a sufficient warrant” (p. 223).<sup>1</sup> Moreover, Mill later calls certain actions ‘self-regarding’ when they do not affect, or at least do not harm, others. For instance, Mill says it is only permissible for society to intervene with an alcoholic if their intemperance leads them “to violate a distinct and assignable obligation to any other person or persons”, since then “the case is taken out of the self-regarding class, and becomes amenable to moral disapprobation in the proper sense of the term” (p. 281). Thus, Mill does employ the notion of self-regarding action, which ought not to be interfered with. It is tempting, therefore, to oppose this to ‘other-regarding’ action, which is potentially liable to social interference. I argue that this traditional reconstruction of Mill’s views is problematic, both in its own terms and as an interpretation of Mill, given his own applications of his principle.

---

<sup>1</sup> All unattributed parenthetical page numbers refer to ‘On Liberty’ in volume XVIII of Mill’s *Collected Works*, edited by J. Robson (Toronto: University of Toronto Press, 1977). References to Mill’s other writings will also be to the *Collected Works* and given by CW, followed by volume and page number.

One preliminary point is that Mill's harm principle actually concerns the *reasons* that may be given for interference, rather than what may be interfered with (Ten, 1980, pp. 40-1; Skorupski, 1989, p. 343). Mill sometimes spoke as if restricting the reasons for interference amounted to restricting what could be interfered with, but this is at best only indirectly so. The harm principle does not tell us what actions should or should not be interfered with, but rather gives us a framework for thinking about intervention. Mill claims that only harm can justify intervention, which excludes certain reasons for intervention, such as preventing harmless immorality. Moreover, not all harms can justify intervention: actions that do not harm *others* fall within a protected sphere of liberty. This protection is not, however, absolute: one cannot intervene with the aim of preventing self-harm, but such actions may nonetheless be interfered with if necessary to prevent other harms (Bird, 2007, pp. 181-6). I shall not dwell on these points, since they have been developed elsewhere (including Saunders, 2013).

One may wish to reject any form of 'harm principle'. For instance, one might argue that the notion of 'harm' cannot be made coherent (Holtug, 2002; Bradley, 2012). Or one may insist that the state should regulate immorality as such, even where it is harmless (Ripstein, 2006). Certainly some generally-accepted legal prohibitions, such as that on incest, are difficult to reconcile with the harm principle.<sup>2</sup> It is not my purpose, here, to respond to all of these objections. I aim to develop the most plausible version of the harm principle but, beyond this, I do not propose to defend

---

<sup>2</sup> Children born of incestuous relationships may be more likely to suffer from genetic abnormalities, but it does not follow that they are harmed as a result, since they would not otherwise exist (the non-identity problem). Moreover, not all incestuous relationships need produce children. However, prohibitions on incest are also problematic for contemporary liberals (Tralau, 2013).

this general approach against those that reject it entirely; in this respect, I have nothing to add to the arguments of *On Liberty*. Commentators have generally focused on the distinction between self-regarding and ‘other-regarding’ actions, with critics alleging that such a distinction is impossible to draw. I shall argue, first, that we can distinguish between actions that are self-regarding and those that are not. Then, more radically, I show that this distinction is not important to Mill or the harm principle. Mill could have dropped this supposedly problematic distinction and still defended all the conclusions he wanted simply by appeal to non-consensual harm.<sup>3</sup> Moreover, this is an independently attractive interpretation of the harm principle. Thus, in thinking about social regulation, we should focus on consent, rather than attempting to delineate a self-regarding sphere.

### **1. Refining the Self-Regarding/Other-Regarding Distinction**

The traditional interpretation of Mill’s harm principle, as relying upon a distinction between self-regarding and other-regarding actions, is not without textual basis. Nonetheless, the distinction is difficult to draw. James Fitzjames Stephen, for instance, famously remarked that “the attempt to distinguish between self-regarding acts and acts which regard others, is like an attempt to distinguish between acts which happen in time and acts which happen in space.... altogether fallacious and

---

<sup>3</sup> In this respect, my reformulation differs from that of Ripstein (2006), since Ripstein proposes his sovereignty principle as an alternative to Mill’s harm principle (albeit suggested by one of Mill’s passing remarks). My proposal is a refinement of the harm principle that, I believe, Mill could and should have accepted. The emphasis that I place on consent appears to move the principle in a Kantian direction (c.f. O’Neill, 1985), but (in keeping with Mill’s utilitarianism) it is still well-being, rather than agency, that is fundamental.

unfounded” (Stephen, 1874, p. x, c.f. pp. 134-50). If everything we do may affect others, then the protected self-regarding sphere will be non-existent.

Sympathetic interpreters have sought to resist this conclusion. John C. Rees suggested that we should focus on actions that affect *the interests of* others (Rees, 1991, pp. 174-80).<sup>4</sup> This interpretation has textual support (pp. 225, 276) but is independently appealing because it preserves some sphere of action that is not ‘other-regarding’ in the relevant sense. Almost anything I do may *causally* affect you, but not everything that causally affects you affects your interests (c.f. Beckman, 2009, pp. 41-7). If we confine our notion of other-regarding action to actions that affect the interests of others, then there will plausibly be a range of individual action that is not other-regarding. This sphere is expanded still further by the observation that not all effects that I might have on your interests license interference; perhaps only harms, i.e. *negative* effects, permit intervention (c.f. Saunders, 2012, p. 290). If I benefit you, that affects your interests, but it is not obviously something that I can be forcibly prevented from doing. (I return to this matter in section 10.) Thus, each of us should be free to perform actions that do not negatively affect the interests of others, either because they do not affect the interests of others *at all* or because they only benefit others.<sup>5</sup>

So interpreted, the harm principle excludes any self-inflicted harm from justifying interference with someone’s liberty. It is, of course, incomplete until we are supplied with an account of people’s interests, without which we are unable to

---

<sup>4</sup> Rees (1991, p. 175) supposes that interests must be socially recognized. I see no reason to assume this.

<sup>5</sup> One question here is whether I am allowed to perform some action that affects you both positively and negatively, provided that on balance you are benefited. There is some *pro tanto* harm here, even if it is outweighed by a greater benefit. I cannot do justice to this issue here.

adjudicate whether or not others have their interests negatively affected. Mill derived his account of interests from his utilitarianism: “I regard utility as the ultimate appeal on all ethical questions [...] Those interests, I contend, authorize the subjection of individual spontaneity to external control, only in respect to those actions of each, which concern the interest of other people” (p. 224). An examination of these permanent interests cannot be attempted here (see Brink, 1992 and Saunders, 2010). Nonetheless, Mill clearly held that human beings in the maturity of their faculties have an interest in liberty. Any action that restricts my liberty, assuming that I am in possession of my faculties, is contrary to my interests and needs justification.

Further, a correct account of interests will doubtless show that there are some things we have no interest in avoiding. Jeremy Waldron (1987) has argued that (in Mill’s view) we have no interest in avoiding distress or offence that comes from moral confrontation. If I express a view contrary to your deeply-held ethical or religious convictions, it may cause you distress, but this does not negatively affect your interests because, on the correct account of interests, you have no interest in avoiding challenges to your convictions. If this is correct, it further expands the sphere of protected liberty. Actions that *seem* harmful to you need not license interference if not actually contrary to your true interests.

We might expand the protected sphere of liberty yet further, for instance by saying that an action is not other-regarding unless it *significantly* affects the interests of others. Mill also allows certain effects on others to be discounted: “When I say only himself, I mean directly, and in the first instance” (p. 225). We might, therefore, refine our notion of self-regarding action, defining it as conduct that does not significantly or, alternatively, directly affect others. These amendments may secure a wider self-regarding sphere, but neither solves all of the problems for the distinction. I

shall return to this point shortly, but for now it will be clearer to stick to a simpler formulation of the distinction.

## 2. Remaining Problems with the Self-/Other-Regarding Distinction

Despite much ingenuity being employed in refining and defending the distinction between ‘self-regarding’ and ‘other-regarding’ action, there remain important difficulties and ambiguities.<sup>6</sup> For one, note that ‘other-regarding’ cannot be understood as the opposite of ‘self-regarding’, at least not if we wish the two categories to be mutually exclusive and exhaustive.

It is natural to assume that ‘other-regarding’ actions are those actions that affect other people’s interests and ‘self-regarding’ actions are those actions that affect the agent’s own interests. But, on these definitions, some actions may be neither self- nor other-regarding, if they do not affect anyone’s interests. A trivial action, like yawning or scratching one’s nose, will not even count as self-regarding, if it does not affect the agent’s interests. The distinction, so construed, does not exhaust all possibilities. More problematically, these two categories are not mutually exclusive. Many actions will be both self-regarding and other-regarding, if they affect the interests of the agent and of others.

Mill’s notion of self-regarding conduct cannot simply be conduct that affects one’s own interests but must mean something like conduct that affects *only* one’s own interests. However, ‘other-regarding’ cannot similarly mean action that affects *only*

---

<sup>6</sup> I shall focus on one particular difficulty, but at least two others are worth mentioning. First, if the consequences of *not* having performed some other-regarding act *Y* are reckoned among the consequences of one’s instead having performed otherwise self-regarding act *X* then *X* becomes other-regarding (Wollheim, 1973, p. 26). Second, Mill presumably did not wish to license paternalistic interventions even in actions that are other-regarding (Saunders, 2013, pp. 73-4).

other people's interests, or else actions that affect both the agent and others will fall between the gaps again, being neither self- nor other-regarding. 'Other-regarding' actions—to fit the sense required—must be any that affect other people, whether or not they also affect the agent. Though the two terms appear opposites, 'other-regarding' cannot be understood analogously to 'self-regarding', which invites confusion. These points hold even on certain refined versions of the distinction. If, for example, we define 'self-regarding' action as what *directly* affects the agent, and other-regarding as what *directly* affects others, then again an action might be both (if it directly affects both the agent and others) or neither (if it does not directly affect anyone).

The categories might be made exclusive and exhaustive if, for instance, we were to define self-regarding as what *primarily* affects the agent, and other-regarding as what *primarily* affects others. This secures a distinction of the sort required, since any action will be either self-regarding or other-regarding, and not both. However, this distinction does not suit Mill's purposes. It is not plausible that society can only intervene with an individual's action when those actions have more effect on others than on the agent herself; surely society may prevent an action that does significant harm to others, even if it has a greater effect (benefit or harm) on the agent.

These problems stem from introducing the term 'other-regarding' to refer to actions that are not self-regarding. Much confusion would have been saved if commentators had never invented the phrase—which, remember, is not Mill's own—and instead spoken of actions being either 'self-regarding' or 'non-self-regarding'.<sup>7</sup> A self-regarding action is one that affects (at least directly) only the agent herself. Non-

---

<sup>7</sup> Just as it is better to contrast instrumental value with *non-instrumental*, rather than intrinsic, value (c.f. Mason, 2001, pp. 247-8).



self-regarding actions are simply those that are not self-regarding, because they (directly) affect others, regardless of any effects they have on the agent herself.

This is our first, rather modest, conclusion. It is unhelpful to introduce the notion of ‘other-regarding’ action, in contrast to self-regarding action. It is clearer and simpler to contrast actions that are self-regarding (directly affecting only the agent) and those that are not, whether or not actions in the latter class are ‘other-regarding’ (that is, directly affecting only others). In saying this, I do not suggest that those who introduced the phrase ‘other-regarding’ are substantively confused about Mill’s position; ‘other-regarding’ is generally used as a term of art to mean what I have called non-self-regarding. My point is a terminological one, but it would be clearer—and more in keeping with Mill’s text—to avoid the potentially misleading phrase ‘other-regarding’.

This goes some way to resolving problems with the notion of self-regarding action. At least a clear line can be drawn, in theory, between the self-regarding and the non-self-regarding; the only question is whether there is any action of consequence on the self-regarding side. However, I now wish to show that this line is not as fundamental to the harm principle as Mill’s interpreters and critics have supposed.

### **3. Expanding the Protected Sphere**

It is unsurprising that many commentators have seized upon the notion of self-regarding conduct, given the importance that Mill himself appears to attach to it. But the ‘self-regarding’ merely marks out one part of the sphere of liberty. Shortly after introducing his harm principle, Mill says “there is a sphere of action in which society, as distinguished from the individual, has, if any, only an indirect interest;

comprehending [*a*] all that portion of a person's life and conduct which affects only himself, or [*b*] if it also affects others, [does so] only with their free, voluntary, and undeceived consent and participation" (p. 225). Mill here identifies *two* parts of the sphere of liberty: first, there is self-regarding action, and, second, actions which affect others only with their consent.<sup>8</sup>

Mill himself sometimes neglects to mention this (e.g. p. 282), but his official view is that consensual harm cannot justify interference (e.g. p. 290), even though the actions that cause it are not self-regarding. Society has no business in protecting people from harms that they do not wish to be protected from, so there is no justification for interference. Thus, for instance, homosexual relationships are protected by the harm principle, though not self-regarding, because the two parties, in consenting to the relationship, consent to any harm that follows.<sup>9</sup> So, even if no action is self-regarding, Mill's principle is not without content: it still says that consensual harm provides no reason for intervention. Excessive focus on the self-regarding/other-regarding distinction can serve to blind us to this point.

---

<sup>8</sup> Mill adds that this consent must be free, voluntary, and undeceived. I shall drop these qualifications, on the assumption that consent secured through duress or deception is not genuine consent. I do not offer a full account of genuine consent; see Miller and Wertheimer (2010).

<sup>9</sup> I assume that the relationship is consensual. Society can intervene in non-consensual relations, whether homosexual or heterosexual. Note also that this seems to assume that consent is transitive, in the sense that if you consent to A, and A entails B, then you consent to B. For instance, if you consent to sexual intercourse, and sexual intercourse carries the risk of pregnancy, then you consent to the risk of pregnancy. This view is criticised by O'Neill (2003, pp. 5-6), who argues that consent, as a propositional attitude, is opaque and so one may consent to A but not B, even though A causes or entails B. Perhaps, however, Mill holds only that consent is transferred to known entailments.

It might be suggested that consensual actions can be understood as ‘self-regarding’ in a wider sense (Athanasoulis, 2002, p. 143). That is, when you do something to me with my consent, or at my request, it might be said the action is not simply yours, but ours: we are acting together (Gilbert, 1989). Indeed, sometimes we may even wish to say that I acted through you.<sup>10</sup> However, even if this is *sometimes* the case, it does not seem plausible that it is always so. There is a significant difference between the case where we jointly do something to me, in which I exercise some agency, and cases where I merely consent, passively, to your doing something to me. The latter are not intuitively self-regarding, unless we use the phrase as a term of art to refer to whatever is protected, in which case action can hardly be said to be protected *because* it is self-regarding: this will be a conclusion, rather than what justifies it. Moreover, whatever *we* think about these cases, Mill clearly distinguishes between two parts of the sphere of liberty, the self-regarding and the consensual (p. 225). We may be uncertain whether some particular acts, such as sado-masochistic sexual practices, are better understood as the self-regarding acts of a collective agent or cases where one person consents to another’s non-self-regarding actions. In practice, it matters little; consensual acts need not be considered joint actions to be protected.

Even if no action is ever self-regarding, one’s action should nonetheless be protected from interference provided that all who are harmed consent. The heavy emphasis on self-regarding action is unfortunate, because it captures only one part of the protected sphere, and perhaps the smaller one at that. Mill’s principle would

---

<sup>10</sup> For instance, suppose that I wish to die but cannot bring myself to do it. I may manipulate you into killing me, yet (depending how the story is filled out) we might still wish to call this a suicide, though I died by your hand rather than my own. I thank Catriona Leyland for related discussions.

appear more promising if greater emphasis were given to the fact that society cannot interfere to prevent consensual harms. (I consider objections to this in section 7.)

#### 4. Narrowing the Protected Sphere

Having shown that Mill's protected sphere includes not only self-regarding action but also what affects others only with their consent, I now wish to suggest that not all self-regarding actions belong to the protected sphere. Many commentators have observed that some of Mill's later remarks in *On Liberty* appear to run contrary to the general anti-paternalism he announces at the start. Two examples are particularly noteworthy. First, Mill suggests that one might hold someone back from an unsafe bridge, long enough to warn him of the danger (p. 294). Second, he insists that slavery contracts should not be enforced (p. 300).

It is often said that these are cases where Mill allowed paternalism (e.g. Garren, 2007, pp. 50-2). To be sure, not all commentators accept that these are genuine instances of paternalism. New (1999, pp. 68-9) distinguishes between interfering with actions based on inadequate information and interventions justified on grounds that the agent interfered with does not reason properly, arguing that only the latter cases constitute genuine paternalism (New 1999, pp. 70-1). Archard (1990, pp. 461-2) denies that the state's refusal to enforce a slavery contract is paternalistic because it is not stopping anyone from living as they wish. These discussions are symptomatic of disagreement about how to understand paternalism (see Garren, 2006). Thankfully, the label 'paternalism' is unimportant here. Both these examples are interventions in self-regarding action (or, at least, justified in order to prevent self-harm) and thus *prima facie* inconsistent with, rather than applications of, the harm

principle as usually interpreted. In each case, the justification that Mill offers for state (in)action rests upon the action's effect on the agent, rather than any harm to others.

In the first case, Mill says that someone can be prevented from crossing an unsafe bridge, though they are the only one at risk of harm. True, this prevention should only be long enough to warn the person in question of the danger; once he is aware of it, Mill insists that he should “not forcibly [be] prevented from exposing himself to it” (p. 294). Note, however, that Mill's position is not simply that we may warn someone of the dangers of their action, but that we may *forcibly restrain* them long enough for us to do so. This involves coercion, and not mere persuasion. Moreover, this coercion is only to prevent the would-be crosser from harming himself, not anyone else. This appears to be a departure from his harm principle and, therefore, to require explanation. The traditional interpretation has it that this intervention can be justified because not contrary to the individual's will (Arneson, 1980, p. 471). This, however, rests upon assumptions about the individual's will and a particular interpretation of what it is to be contrary to someone's will.<sup>11</sup> Furthermore, it complicates the self-regarding/non-self-regarding distinction by adding a separate condition of voluntariness. The alternative that I propose, below, is simpler and more unified.

---

<sup>11</sup> It is assumed that if the individual has no will then nothing can be contrary (or opposed) to her will. A broader understanding would have it that anything not in accordance with one's will is contrary to one's will; thus intervention may be contrary to one's will even where one has no particular will. The former, narrower reading may be preferable, but is not the only possibility.

Moreover, even if the individual does not wish to fall in the river, it is not clear that he would consent to being interfered with. Someone very averse to outside interference might reject it, even if it meant falling in the river. For problems with hypothetical consent, see Hanna (2012, pp. 428-32).

The second case in which Mill allows for apparently paternalistic intervention concerns slavery contracts. Mill does not actually say that such contracts should be prohibited, only that they should not be enforced (p. 299; c.f. Archard, 1990, pp. 461-2). However, presumably this means not only that such contracts will not enjoy the usual protection of law, but also that parties should not be able to arrange their own private enforcement as part of the contract. Assuming no party will sign up to a contract where they cannot be assured of the other's compliance, such contracts are rendered effectively impossible. Mill suggests that such contracts are not a permissible exercise of liberty, because they forfeit the very liberty that must be appealed to in order to protect them (p. 300). It seems that Mill thinks it something like a performative contradiction to renounce one's freedom, since one can only justify one's right to do so by appealing to the liberty that one wishes to renounce (Lovett, 2008, p. 130; c.f. Archard, 1990, p. 459).

It might be replied that prohibiting slavery contracts is not in tension with Mill's harm principle, because a contract is not self-regarding: it requires another party. To be sure, Mill does say that "trade is a social act [and] comes within the jurisdiction of society" (p. 293). Slavery contracts need not involve exchange though; I might offer to enslave myself to you as a gift, asking nothing in return. Still, even gifts can be prohibited if they are harmful to their recipients. Mill suggests that slavery debases the master as well as the slave (p. 269; see also *CW XXI*, p. 325), so he could have rejected slavery on these grounds. However, while Mill *could* consistently have argued against slavery by appeal to the interests of others, the arguments he *actually* offered concern the effect on the would-be slave. Even if the action is not strictly self-regarding, the reasons offered for intervention do not appeal

to harm to others but the effect on the agent herself.<sup>12</sup> The mere fact that another's interests are involved should not license appeal to paternalistic reasons for intervention (c.f. Saunders, 2013, pp. 73-4).

The harm principle supposedly tells us that we can never interfere with an individual's conduct except to prevent harm to others. However, Mill sometimes allows interventions to protect the individual from her own action. Thus, even if the self-regarding/non-self-regarding distinction is salvageable, it does not fit Mill's own applications of his harm principle. I will now propose an alternative formulation of Mill's harm principle, which better captures his intentions.

### **5. The Principle Reformulated**

I suggest that we reformulate Mill's harm principle as follows. Instead of saying "the only purpose for which power can be rightfully exercised over any member of a civilized community, against his will, is to prevent harm to others" (p. 223), what Mill *should* have said is that the only justification for interfering in someone's liberty is to prevent *non-consensual harm*. This has the advantage of making plain that there is no basis for interfering in conduct that only harms consenting others, while allowing interference in self-regarding conduct where the agent does not consent to the harm she does herself.

I believe that this is in keeping with Mill's intentions, since the two cases of paternalism discussed above can both be understood as ones in which the agent

---

<sup>12</sup> A similar response also applies to those, like Hodson (1981) and Brown (1989), who argue that non-enforcement does not diminish liberty. Mill may have been wrong to think that a refusal to enforce slavery contracts denies people liberty, but he was still prepared to limit liberty.

cannot consent to harm (because unaware of it<sup>13</sup> and because, in Mill's view, logically incoherent). While Mill distinguishes *two* separate spheres of liberty, self-regarding conduct and what harms others only with their consent, I suggest that the protected part of the former category could be subsumed within the latter. What really matters, for Mill, is whether those harmed by the conduct in question consent to that harm.


On this interpretation, the sphere of protected action is not the strange disjunct of actions that are *either* self-regarding or consented to, but rather self-regarding actions in which the individual consents to any harm and non-self-regarding actions in which the individuals concerned consent to any harm. In both cases, what does the normative work is the consent of those harmed. The distinction between self-regarding and non-self-regarding action is irrelevant, except that it marks the need for consent from others besides the agent herself. We can represent this diagrammatically:

---

<sup>13</sup> There is a complication here if the individual is culpably ignorant (Hanna, 2012, p. 424). Perhaps we may say his consent is not valid but, because he is culpable for its invalidation, we may treat him *as if* he had validly consented, and thus cannot interfere. For the idea that non-consent can be rendered void, though not exactly as described here, see Estlund (2008, pp. 121-7).



Fig. 1:



This circle represents the range of actions that one may perform. The line A-B is the distinction between self-regarding actions (sections 1 and 2) and non-self-regarding actions (sections 3 and 4). The line C-D is the distinction between consensual actions (sections 1 and 3) and non-consensual actions (sections 2 and 4).

Emphasizing line A-B neglects the fact that Mill's sphere of liberty explicitly includes section 3, i.e. actions that affect others only with their consent (pp. 225-6). While there is no inconsistency in allowing *either* self-regardingness *or* consent to protect liberty, this disjunctive account is less unified than one in which a single feature distinguishes the protected sphere. Furthermore, as we have seen, Mill *is* prepared to countenance some interference in self-regarding actions. These exceptions can be explained by the fact that the actors in question do not consent to the harm they do themselves, which is to say that section 2 here is not protected. This category may

be small, which may explain why Mill neglects it, but it includes cases where the agent harms themselves through ignorance (as in the bridge-crossing case) or temporary incapacity (for instance, harmful choices made under the influence of alcohol or drugs).

If section 3 (non-self-regarding, consensual action) is protected liberty, and section 2 (self-regarding, non-consensual action) is not, then the line that demarcates the protected sphere is actually C-D and not A-B (or, for that matter, C-B). Thus, Mill's harm principle does not rest on a distinction between the self-regarding and non-self-regarding, but says that society has no business interfering in conduct except to prevent non-consensual harm. Put this way, I believe the principle is intuitively clear, plausible, and better conforms with Mill's own applications. It does not matter whether an individual harms herself or others, provided that all who are harmed consent to the harm. This is not to say that such actions can never be interfered with, if we understand the harm principle as regulating reasons for intervention (Ten, 1980, pp. 40-1). But, whereas the traditional interpretation has it that self-harm is never an acceptable reason for intervention, this interpretation focuses instead on consent. Consensual harm is never grounds for intervention, while intervention to prevent self-harm can be justified where that harm is non-consensual.

Ordinarily, self-regarding actions will be protected, because people generally consent to any harm that they knowingly do to themselves.<sup>14</sup> What explains the

---

<sup>14</sup> Feinberg (1971, p. 107) regards talk of consenting to one's own actions as metaphorical. But if the notion of an individual consenting to her own action seems odd, it is presumably because we ordinarily assume that individuals willingly accept the (foreseen) consequences of their own actions. We rarely talk of consent except in more formal contexts, but this does not make it inapplicable (O'Neill, 2003, p. 4). While this language might be unusual, I believe it fruitful because it brings out a parallel between the self-regarding and non-self-regarding cases.

immunity of these actions from interference, however, is *not* simply that they are self-regarding, but that the person harmed (the agent) consents. The bridge-crosser, conversely, does not consent to the risk of falling in the river, because he is unaware of it. And the person who would subject himself to slavery cannot coherently do so, Mill suggests, because it involves some form of contradiction. These actions are not protected, because the harm they involve is non-consensual.

This reasoning can explain our reaction to cases that Mill does not explicitly consider. The requirement that others be (prejudicially) affected, if at all, “only with their free, voluntary, and undeceived consent and participation” (p. 225) excludes, for instance, drunken consent. Mill is clear not only that children lack the power to consent but also that a person’s liberty may be interfered with if he is “delirious, or in some such state of excitement or absorption incompatible with the full use of the reflecting faculty” (p. 294). If someone is too drunk to consent, then they cannot give their consent to, e.g., sexual intercourse or to getting a tattoo. It would, ordinarily, be a violation of individual liberty to prevent people from engaging in sexual intercourse or getting tattoos, though these actions are not strictly self-regarding (they involve other people doing things to you); but it is not a violation of one’s liberty if one is unable to consent. Again, the law serves to protect people from harms that must be non-consensual, because people cannot give valid consent in such conditions. These are not *ad hoc* restrictions of the principle’s scope, but implications of it.

One might object that consensual harm is an oxymoron. If we accept the principle *volenti non fit injuria* (to a willing person, no injury is done), then we may think that someone who consents to something is not injured by it (c.f. Lovett, 2008, p. 128). This objection, however, conflates the notions of harm (a setback of interests) and injury (a wrongful or unjust harm) (Lazar, 2009). One who consents to something

is not wronged by it, but may still be harmed (c.f. Feinberg, 1971, p. 108). Some commentators have attributed a ‘moralized’ notion of harm to Mill (see Holtug, 2002, 377ff.), but it seems that Mill thought of harm in the descriptive sense, since he allows that individuals may be harmed in permissible ways, such as through competition (pp. 292-3). Moreover, we can surely be harmed, though not wronged, by things that are not moral agents, such as animals or natural disasters (Bradley, 2012, pp. 394-5). Thus, we can, and Mill could, coherently distinguish between consensual and non-consensual harms. There is no reason to suppose the former particularly problematic, once we recognize that harms need not be wrongful and consent merely negates wrongness, and not harm.

A distinction between consensual and non-consensual harm better captures the intentions behind the harm principle. It explains both why we *cannot* intervene in some non-self-regarding actions (because the harm is consensual) and why we *can* sometimes intervene in some self-regarding actions (because the harm the individual does to himself is not consensual). Thus, the distinction between self- and non-self-regarding conduct is unnecessary; Mill could simply have referred to non-consensual harm, instead of harm to others. To be sure, this is not what he actually said. Nonetheless, I believe that this reformulation is substantively plausible and better captures Mill’s own views.

## **6. Was this Mill’s View?**

The reformulation offered in the last section is intended not simply as a replacement for Mill’s harm principle, but as an alternative statement that he could have accepted. The primary justification for attributing this to Mill is that it fits with his various examples and applications, including cases of conduct that can be

interfered with to prevent agents from harming themselves. It might, however, be objected that we ought to focus on Mill's official statements of his principle, rather than his examples, regarding the latter as possible misapplications. Of course, where there are apparent inconsistencies within a text, it is difficult to say for certain which passages capture the author's true, considered opinion and which may be dismissed as rhetorical exaggeration or a slip of the pen. However, I think we have reason to take Mill's initial statement of the principle as an exaggeration and his later examples as better indicating his real, more nuanced views.

While it is not conclusive, some support can be found in Mill's essay 'Bentham', where he wrote that "all writing which undertakes to make men feel truths as well as see them, does take up one point at a time, does seek to impress that, to drive that home... It is justified in doing so, if the portion of truth which it thus enforces be that which is called for by the occasion. All writing addressed to the feelings has a natural tendency to exaggeration... we must aim at too much, to be assured of doing enough" (*CW X*, p. 114). This passage is followed by criticism of Bentham's style, in which Mill attributes Bentham's "intricate and involved style" to an "impracticable precision" which insists on incorporating all intended qualifications into every sentence, rather than being prepared to say "a little more than the truth in one sentence, and correct it in the next" (*CW X*, p. 114). Given that Mill sought to write for a general audience, it is reasonable to assume that he was willing to do what Bentham was not.

In light of these remarks, it is plausible to postulate that *On Liberty* sometimes overstates Mill's views, in order to counterbalance what he regarded as a tendency towards excessive social interference in individual liberty (pp. 219-20). Faced with this danger, he stated the case for individual freedom in more absolutist terms than

would otherwise be warranted. To be sure, Mill did not (so far as I am aware) admit to employing such rhetorical strategies in *On Liberty*. However, he did say that this is necessary in *all* writing intended to drive a truth home and surely *On Liberty* had this intention. Therefore, I believe that the best way to resolve the apparent inconsistencies between Mill's initial statements of his principle and his applications of it is to suppose that he was never really an absolutist, even about self-regarding liberty, but overstated his case because he felt that the social tendencies of his time were running the other way.

I believe that this position is more plausible than the one traditionally attributed to Mill. In the next section, I consider whether excluding interference in all consensual harm is too restrictive; then, in section 8, I will consider whether this reformulated principle permits too much paternalistic interference.

### **7. Should All Consensual Harm be Protected?**

I have argued that consensual harm is not grounds for interference. One might ask whether this allows scope for criminal law. We do not ordinarily believe that consent makes all harms permissible; there are some gross harms (maims) that we think it the business of society to prevent, even where the victim consents (Feinberg, 1971, pp. 105-6; Baker, 2009, pp. 97-8). If the harm principle prohibits society from interfering whenever harms are consented to, then it may prohibit us from interfering in these cases. It should be noted, however, that this is an objection not only to my reformulation of the harm principle, but for Mill's official statement too, since he includes consensual harm within the protected sphere (p. 225). This objection may lead us to reject the harm principle altogether, but does not tell against my claim that

formulating it in terms only of consent is the simplest and most plausible formulation of it. Even so, Mill has resources to respond to such objections.

First, it should be stressed that the purpose of the principle is precisely to limit the ability of the state and society to interfere in individual freedom. Perhaps we *want* to interfere with consensual sado-masochism, but it does not follow that we have the *right* to, any more than consensual homosexuality (c.f. Athanassoulis, 2002). However, if we wish to accommodate the view that the state ought to prevent certain harms, even where those harmed appear to consent, we might begin by noting that the consent involved must be genuine. We may worry whether apparent consent is genuine in some cases (Arneson, 1980, p. 472). We should not exclude the possibility of genuine consent *a priori*, but there are certainly cases where, on hearing that someone has apparently consented to some harm, it is more reasonable to assume that this ‘consent’ is defective (e.g. they are coerced or not thinking clearly) than that the consent is genuine. In these cases, interference may be permissible, at least until it can be ascertained that the consent really is genuine.

Further, even if the most immediate ‘victim’ of some act consents to the harm, that does not show that the act cannot be interfered with to protect others from being harmed without their consent. For instance, even if the ‘stars’ of violent pornography consent to their role in its production, society might intervene if it is harmful to other women who do not consent (c.f. Dyzenhaus, 1992; Skipper, 1993; Vernon, 1996).

Additionally, that society may only interfere with an individual’s actions in order to prevent (non-consensual) harm does *not* imply that only (non-consensual) harm-causing actions may be interfered with. Law, and even social opinion, is a blunt instrument. Sometimes it is necessary to interfere with harmless actions in order to prevent harmful ones (Bird, 2007, pp. 181-6). While a more targeted intervention is,

*ceteris paribus*, to be preferred (Saunders 2013, pp. 76-7), where that is not possible a broader intervention may be justifiable, provided the harm prevented is great enough to warrant a larger loss of liberty. So, even if some people genuinely consent to harm *H*, it may be that, given difficulties in determining whether consent is genuine, the only way to protect others from suffering *H* without consent is to prohibit any action of inflicting *H*. Thus, even if some people do genuinely consent to maims, the state may be justified in intervening in all cases, because such consent is often suspect and it cannot be sure which cases are genuine.

I believe that Mill need not be embarrassed by the claim that consensual harm cannot justify interference. However, even if these arguments are not convincing, it should be remembered that these cases are not problematic for my reformulation of the harm principle in particular, since Mill's official statement also excludes consensual harm from grounds for interference. My purpose has been to defend the substantive plausibility of this position, but where my reformulation differs from Mill's is that it allows for more interference.

### **8. Is (Soft) Paternalism Permissible?**

The substantive difference between my reformulation of the harm principle and Mill's original statement, which includes consensual harm in the protected sphere (p. 225), is that mine explicitly allows for interference in self-regarding conduct where the individual does not consent to the harm that she does to herself. That is, this formulation explicitly allows 'soft' paternalism, which seems to be precluded by some of Mill's absolutist statements (even if, as we have seen, he allows some such interventions in his applications).



I believe that a principle that allows for the possibility of justified soft paternalism is more plausible than one that absolutely excludes it. I cannot defend the permissibility of soft paternalism here, but if soft paternalism is even *potentially* justifiable, it seems better that the principle allows for the possibility. A principle that potentially permits such intervention, by not placing non-consensual self-harm in the protected sphere, does not mandate paternalistic intervention: it is still possible to reject soft paternalism in each individual case. Thus, the reformulated principle is consistent with soft paternalistic interventions occurring or not occurring, whereas a version that absolutely forbids intervention to prevent self-harm prejudices the issue. The reformulated principle is not only more consistent with Mill's applications (some of which involve soft paternalism), but independently more plausible, because more permissive.

It may be objected that this permissiveness comes at the cost of indeterminacy: the principle does not settle whether or not a given paternalistic intervention is justified. However, the principle does not directly settle the justifiability of any intervention; it merely frames how we should think about justification, by settling what reasons we can and cannot appeal to (p. 292). In debating whether or not a particular intervention is justified, we will be called upon to consider various factors, including balancing the harm prevented against the loss of liberty. For Mill, this calculation consists in comparing utilities (c.f. Brown, 2010), though we need not accept this particular detail. Thus, the harm principle itself does not usually settle the justifiability of particular interventions, except where there is no permissible reason for interference (though, even here, see the arguments of the previous section), but merely tells us *how* to go about settling them. This is true for both Mill's original formulation and my reformulation. Both tell us that non-consensual harm cannot

justify intervention; the difference is that my reformulation allows for (but does not require) interventions designed to prevent individuals from harming themselves, where they do not consent to that harm. This, I believe, is a more attractive principle, though whether or not we should intervene in such cases will depend on the costs and benefits of doing so in a particular situation.

Those who advocate a particular intervention need not do so in terms of harm to others, which may not be their real motive for interference, but are free to appeal to self-harm, provided they can show that agents do not freely consent to this harm. Thus, in cases like alcohol (see Saunders, 2013) or unhealthy food (see Resnik, 2010), they may argue that many people do not properly appreciate the long-term health costs of their choices. They may add that people's choices are distorted by poor nutritional information, advertising, and social pressures (Gostin, 2010; Kirkwood, 2010). Mill was also concerned with such influences (pp. 296-7), which is readily explained by a focus on genuine consent. While the presumption should always be in favour of liberty, and against interference, placing the onus on would-be interferers to show that consent is not genuine, we need not suppose that self-harm is never reason to interfere.

## **9. Further Questions**

Reformulating the harm principle in terms of non-consensual harms still leaves unanswered a number of important questions. We still need an account of what interests people have, though this is true also for the traditional version of the principle. We also need some account of valid consent. Mill notes that consent must be free, voluntary, and undeceived. We might add, in light of the unsafe bridge case, that it should be informed: one cannot consent to the harms if one is unaware of

them.<sup>15</sup> Further, one also needs the *capacity* to consent. We can take Mill's restriction of the principle to those "in the maturity of their faculties" (p. 224) not as an *ad hoc* scope restriction of the principle, but rather an implication of it. If minors cannot give valid consent, then harms to them will always be non-consensual.

There are many important questions to resolve here, for instance, how much information someone needs to have for their consent to be 'informed'. Presumably we do not require that people possess *all* possibly relevant facts, so 'informed' must mean something like possessing (or having access to) the information that a reasonable, prudent decision-maker would want (O'Neill, 2003, p. 6). Another issue is that decision-making competence is not all-or-nothing but a matter of degree. A fourteen-year-old may be competent to make some decisions over her own life, like what to eat, but not competent to make more momentous decisions, including perhaps whether to have sexual intercourse. Recognizing this is, I submit, an improvement over any view that suggests a single 'age of maturity'. Further, what we need, as satisfactory proof of consent, may also differ depending on what it is that is being consented to. If the possible harm is minimal or unlikely, then the standards of consent that we demand may be low, but we may need correspondingly better evidence where the risk of harm is greater.

To provide a full account of valid consent, either in fact or in Mill's view, is beyond the scope of the present article. Note that its necessity is not a result of my reformulation of the harm principle; it will be required anyway, since Mill includes consensual harm, as well as self-regarding harm, in the protected sphere. My reformulation at least has the advantage of making these the only problems, showing

---

<sup>15</sup> Mill introduces the unsafe bridge case in order to defend the labelling of dangerous drugs, because "the buyer cannot wish not to know that the thing he possesses has poisonous qualities" (p. 294).

that there is no need to concern ourselves also with trying to distinguish between actions that are self-regarding and actions that are not. If we want to think about social regulation in a Millian way, it is more fruitful to focus our attention on working out the details of consensual harm (thus, this account provides a future research programme, providing an account of valid consent, as well as suggesting consent as a focus for practical deliberation).

### **10. Harmed or Affected?**

One further matter deserves comment. I have focused on non-consensual *harms*. These are, first, non-consensual, and, second, harmful. I assumed that society has no jurisdiction over conduct that *benefits* others. We might question this, saying that individuals have no more right to benefit others unilaterally than they do to harm them unilaterally. Ordinarily, we may expect that individuals will readily consent to benefits, but, if this consent is not forthcoming, perhaps society may intervene.<sup>16</sup> We may therefore distinguish two different versions of my reformulated harm principle:

*Narrow version:* Society may intervene only to prevent actions that *harm* others without their consent.

*Wide version:* Society may intervene only to prevent actions that *affect* others' interests without their consent.<sup>17</sup>

---

<sup>16</sup> Presumably one reason why it may not be forthcoming is that the individual considers the benefit to come with a cost, such as a sense of indebtedness. These cases are instances of *pro tanto* benefits accompanied with harms. See footnote 5.

<sup>17</sup> I call this version 'wide' because it permits a wider range of social interference: society is permitted to intervene to prevent both non-consensual harms and non-consensual benefits.

The narrow version assumes an asymmetry between benefits and harms. On this version of the principle, I should be free to act in ways that benefit you, even if you (wrongly) regard these benefits as contrary to your interests.<sup>18</sup> On the other hand, I can act in ways that are objectively harmful to you, provided that you consent, because you (wrongly) think them beneficial. Thus, effects on others are licensed if objectively beneficial, even if the individual in question regards them as harms, but objective harms need not license interference, if the individual in question does not regard them as such. It seems puzzling that the individual should have discretion over harms, but not over benefits.

The wide version, conversely, gives individuals discretion over whatever affects their interests. Not only can they consent to be harmed, but they can refuse benefits. We might find this more appealing, since it treats benefits and harms symmetrically and allows individuals more control over their own lives: if you wish to impose a benefit upon me, then I can reject it. This does not mean that society *must* prevent you from imposing this benefit, but that it could be legitimate in doing so. Ordinarily we may suppose that individuals will consent to be benefited, and that it

---

<sup>18</sup> While one of Mill's arguments for the harm principle is that individuals are, in general, more likely to know their own interests than others are, his contrast between Socrates and the fool in chapter 2 of 'Utilitarianism' (CW X, p. 212) makes evident that individuals are not *infallible* judges of their own interests. I believe that Mill operates with an objective notion of interests (see Saunders, 2010 and 2011), but the argument does not depend on this. Suppose that Mill operates with a subjective notion of interests, so whatever you think is good (bad) for you is good (bad) for you. You may think that having me do *X* to you would be good for you, but not want me to do *X* to you (perhaps because you do not want to feel indebted to me). What you think good for you is may be different from what you want to happen; you may consent to something, though it is not good for you, or refuse consent to something that is good for you. A subjectivism that cannot make sense of these possibilities is implausible.

would not be expedient for society to prevent non-consensual benefits, but there might be cases in which society may wish to prevent someone from imposing benefits on unwilling recipients.<sup>19</sup> On the wide version of the harm principle (which, despite the name, focuses on *any* non-consensual affects, and not merely harm), this could be legitimate.

I shall not attempt to adjudicate, here, which of these two versions of reformulated harm principle Mill would, or we should, accept. For now, I simply wish to highlight the distinction between these two versions of the ‘harm’ principle and their differing implications.

## 11. Conclusion

It is often supposed that Mill’s harm principle rests upon a distinction between self-regarding actions, which are immune to interference, and other-regarding actions, which are liable to interference. As many commentators have observed, this distinction is problematic. I have argued that a distinction can be drawn between self-regarding and non-self-regarding actions. However, this does not capture Mill’s meaning, since he rightly allows intervention in some self-regarding cases, while prohibiting it in many non-self-regarding cases (where consent is present). That so many interpreters have focused on the self-regarding/other-regarding distinction has blinded many to the fact that it is neither crucial to Mill’s position nor the most plausible formulation of the harm principle.

---

<sup>19</sup> Discussion provides an interesting test case. Waldron holds that a progressive being has an interest in having her convictions challenged, whether she likes it or not. If she is entitled to reject this benefit, then censorship may be permissible. Note, however, that Mill describes silencing an opinion as “robbing the human race” (p. 229). Perhaps it would only be legitimate if *everyone* rejected the benefit (i.e. one individual cannot refuse a benefit if they thereby deprive others who may want it).

I suggested that Mill's intentions would have been better served if he had said that the only legitimate ground for interfering in someone's liberty is to prevent non-consensual harm (or, perhaps, non-consensual effects). This explains why we are usually prohibited from intervening with self-regarding conduct, because individuals typically consent to what they do to themselves, while also explaining exceptions to this rule: it is permissible to stop individuals from harming themselves in ways that they do not consent to. This reformulation emphasizes that other-regarding conduct is equally sacrosanct where those harmed consent to that harm. Thus, the crucial test for the legitimacy of social intervention is not where harm falls, but whether it is consented to by those that suffer it. While this reformulation is not necessarily unproblematic, all of its problems are shared with Mill's official position, while those problems associated with the self-regarding/non-self-regarding distinction are rendered nugatory. The permissibility of intervention to prevent harm should depend simply on whether those harmed consent to that harm, rather than who is harmed.

#### Acknowledgements

Previous versions of this paper were presented at the University of Edinburgh (January 2012), the Joint Sessions conference at the University of Stirling (July 2012), and Balliol College, Oxford (March 2013). I thank these audiences, particularly Liz Cripps, Antony Duff, Chris Macleod, James Mahon, Mark Philp, Jeremy Waldron, and Jo Wolff, for their discussions. Particular thanks go to Simon Hope and three anonymous referees for detailed written comments.

#### References

Athanassoulis, Nafsika. (2002) 'The Role of Consent in Sado-Masochistic Practices' *Res Publica* 8: 141-55.

Archard, David. (1990) 'Freedom Not to be Free: The Case of the Slavery Contract in J. S. Mill's *On Liberty*' *Philosophical Quarterly* 40: 453-465.

Arneson, Richard. (1980) 'Mill versus Paternalism' *Ethics* 90: 470-89.

Baker, Dennis. (2009) 'The Moral Limits of Consent as a Defense in the Criminal Law' *New Criminal Law Review* 12: 93-121.

Beckman, Ludvig. (2009) *The Frontiers of Democracy: The Right to Vote and its Limits*. Basingstoke: Palgrave-Macmillan.

Bird, Colin. (2007) 'Harm versus Sovereignty: A Reply to Ripstein' *Philosophy & Public Affairs* 35: 179-194.

Bradley, Ben. (2012) 'Doing Away with Harm' *Philosophy & Phenomenological Research* 85: 390-412.

Brink, David. (1992) 'Mill's Deliberative Utilitarianism' *Philosophy & Public Affairs* 21: 67-103.

Brown, Don G. (1989) 'More on Self-Enslavement and Paternalism in Mill' *Utilitas* 1: 144-150.

Brown, Don G. (2010) 'Mill's Moral Theory: Ongoing Revisionism' *Politics Philosophy & Economics* 9: 5-46.

Burley, Justine. (2008) 'Mill, Liberty, and (Genetic) 'Experiments in Living'' in C. L. Ten (ed.) *Mill's On Liberty: A Critical Guide* (Cambridge: Cambridge University Press), pp. 185–208.

Dworkin, Gerald. (1984) 'Paternalism: Some Second Thoughts' in Ralf Sartorius (ed.) *Paternalism* (Minneapolis: University of Minnesota Press), pp. 105-111.



Dyzenhaus, David. (1992) 'John Stuart Mill and the Harm of Pornography' *Ethics* 102: 534-551.

Estlund, David. (2008) *Democratic Authority: A Philosophical Framework* (Princeton: Princeton University Press).

Feinberg, Joel. (1971) 'Legal Paternalism' *Canadian Journal of Philosophy* 1: 105-26.

Garren, David. (2006) 'Paternalism, Part I' *Philosophical Books* 47: 334-41.

Garren, David. (2007) 'Paternalism, Part II' *Philosophical Books* 48: 50-9.

Gilbert, Margaret. (1989) *On Social Facts* (Princeton: Princeton University Press).

Gostin, Lawrence. (2010) 'Trans Fat Bans and the Human Freedom: A Refutation' *American Journal of Bioethics* 10(3): 33-34.

Hanna, Jason. (2012) 'Paternalism and the Ill-Informed Agent' *Journal of Ethics* 16: 421-39.

Hodson, John. (1981) 'Mill, Paternalism, and Slavery' *Analysis* 41: 60-62.

Holtug, Nils. (2001) 'The Harm Principle and Genetically Modified Food' *Journal of Agricultural and Environmental Ethics* 14: 169-178.

Holtug, Nils. (2002) 'The Harm Principle' *Ethical Theory and Moral Practice* 5: 357-389.

Kirkwood, Kenneth. (2010) 'Lipids, Liberty, and the Integrity of Free Actions' *American Journal of Bioethics* 10(3): 45-46.

Lazar, Seth. (2009) 'The Nature and Disvalue of Injury' *Res Publica* 15: 289-304.

Lovett, Frank. (2008) 'Mill on Consensual Domination' in C. L. Ten (ed.) *Mill's On Liberty: A Critical Guide* (Cambridge: Cambridge University Press), pp. 123-37.

Mason, Andrew. (2001) 'Egalitarianism and the Levelling Down Objection' *Analysis* 63: 246-254.

Mill, John Stuart. (1963-91) *Collected Works*, edited by John M. Robson (Toronto: University of Toronto Press).

Miller, Franklin and Alan Wertheimer. (eds.) (2010) *The Ethics of Consent: Theory and Practice* (New York: Oxford University Press).

New, Bill. (1999) 'Paternalism and Public Policy' *Economics and Philosophy* 15: 63-83.

O'Neill, Onora. (1985) 'Between Consenting Adults' *Philosophy & Public Affairs* 14: 252-277.

O'Neill, Onora. (2003) 'Some Limits of Informed Consent' *Journal of Medical Ethics* 29: 4-7.

Rees, John. (1991) 'A Re-Reading of Mill On Liberty' in J. Gray and G. W. Smith (eds.) *J. S. Mill On Liberty in Focus* (London: Routledge), pp. 169-89.

Resnik, David. (2010) 'Trans Fat Bans and Human Freedom' *American Journal of Bioethics* 10(3): 27-32.

Riley, Jonathan. (1991) 'One Very Simple Principle' *Utilitas* 3: 1-35.

Ripstein, Arthur. (2006) 'Beyond the Harm Principle' *Philosophy & Public Affairs* 34: 215-245.

Saunders, Ben. (2010) 'J. S. Mill's Conception of Utility' *Utilitas* 22: 52-69.

Saunders, Ben. (2011) 'Re-interpreting the Form of Qualitative Hedonism Advanced by J.S. Mill' *The Journal of Value Inquiry* 45: 187-201.

Saunders, Ben. (2012) 'Defining the Demos' *Politics, Philosophy and Economics* 11: 280-301.

Saunders, Ben. (2013) 'Minimum Pricing for Alcohol: A Millian Perspective' *Contemporary Social Science* 8: 71-82.

Skipper, Robert. (1993) 'Mill and Pornography' *Ethics* 103: 726-730.

Skoruspki, John. (1989) *John Stuart Mill* (London: Routledge).

Stephen, James F. (1874) *Liberty, Equality, Fraternity* second edition (London: H. Elder & Co.).

Ten, Chin L. (1980) *Mill on Liberty* (Oxford: Clarendon).

Tralau, Johan. (2013) 'Incest and Liberal Neutrality' *Journal of Political Philosophy* 21: 87-105.

Vernon, Richard. (1996) 'John Stuart Mill and Pornography: Beyond the Harm Principle' *Ethics* 106: 621-632.

Waldron, Jeremy. (1987) 'Mill and the Value of Moral Distress' *Political Studies* 35: 410-423.

Wollheim, Richard. (1973) 'John Stuart Mill and the Limits of State Action' *Social Research* 40: 1-30.