

Unpleasant Realities and the Prudential Good

Valerie Tiberius
<tiberius@umn.edu>

This is a very rough draft: comments are welcome, but please do not cite.

If you were seriously self-deceived, unaware of what others really think of you and what your talents really are, would you want to know? If your significant other had an affair, do you think you'd be better off knowing? If you were living in a pleasant, but – unbeknownst to you – fictional and televised reality (as in “The Truman Show”) do you think it would be good for you to get out of it? How important is it to know the truth about ourselves, our relationships, and our circumstances? What is the relationship between facing up to unpleasant realities and our own good? Certainly facing the facts has vital instrumental value – we will usually make better decisions if we know the facts. But the question I want to focus on here is whether facing unpleasant truths is something we should value for its own sake. Is it good for us to be in touch with reality whether or not it pays off in other terms?

A number of philosophers have thought so. Nozick raised the point as an objection to hedonism. He argues that certain things matter to us other than our experiences: “we want to *do* certain things, and not just have the experience of doing them... we want to *be* a certain way...” and we want to be open to contact with a deeper reality (1974: 43). There's also a long tradition in moral philosophy (from Natural Law Theory to Moore's intuitionism) according to which knowledge is an intrinsic good.

Others are less sanguine about the value of facing the facts. Many philosophers who defend subjective theories of the good for a person tend to add qualifiers that aim to capture the intuition that we're better off when we are in touch with reality, but these qualifiers are often dependent on how people themselves feel about having the information in question. For example, according to Sumner's (1996) life-satisfaction account of well-being, a person's well-being is discounted only if the person herself would not be satisfied if she had better information. In his defense of hedonism, Feldman (2004) provides the option of “truth-adjusted” pleasures for those who think that pleasures in the experience machine count less than pleasures gained in the real world. Neither thinks that being informed is a good thing, period. It's good for you if it matters to you, but not otherwise.

Lack of consensus invites the question: How should we decide whether facing the facts is good for its own sake? In philosophy (and perhaps also in practice), the method for investigating the relationship between well-being and knowing the truth has been, by and large, to consult our intuitions about the nature of a good life for a human being. Are intuitions a good guide? Consider this anecdote, reported to me by a friend who is a professional philosopher:

My husband discovered that a friend of ours had cheated on his wife (also a friend). The affair was brief, unlikely to be repeated, and did not mean that there was anything wrong with the friends' marriage. So he asks me whether I should tell the wife. “No, definitely not”, I said, with great confidence. We talked some more and agreed that this was the

right answer. “So”, my husband asks next, “does that mean if I were the one who had an affair you wouldn’t want to know?” “Of course I would!!”, I said, again very confidently, but surprised by my inconsistency.

My friend had different intuitions about the importance of knowing the truth depending on whether she was thinking about her own case or the case of her friend. In conversation with other people, I have found that this combination of intuitions is not at all uncommon. If we do have different intuitions about the question “Is it good to face the facts?”, which ones should we consult? Unless there is some important difference between oneself and the other person, it can’t be that both intuitions are correct. What would this mean for our question and for methods of investigating well-being more generally?

Of course, a few anecdotes and conversations do not suffice to reveal a general problem with our intuitions. Before we begin considering the philosophical implications of a systematic bias in our intuitions we need to know whether there really is such a bias. I propose to investigate this question experimentally.¹ Toward this aim, I have conducted a pilot study in which I asked people about the relationship between knowing the truth and well-being. What I plan to do here is to report what I’ve found already, to discuss the philosophical implications of the results, and then to raise a number of questions about how to go forward.

Before we begin, it will be helpful to address at the outset one skeptical response to this proposed line of research. One might argue that consulting intuitions is the wrong way to go about answering questions about well-being in the first place, and that what we really ought to do is to consult our *theory* of well-being. If you have an objective theory of well-being according to which knowledge is intrinsically good, facing the facts will be good for a person unless it is outweighed by other intrinsic goods. On the other hand, if you have a desire based theory, whether it is good to face the facts will depend on whether you want to face them, or whether some ideal version of you would want you to want to face them. This criticism fails to acknowledge the ways in which intuitions about well-being inform the theories that we have. Nozick’s case against hedonism, after all, relies on our intuitions about life in the experience machine. Desire based theories are attractive, in part, because of the intuition that if you get some objective good but you don’t want it as part of your life, then it doesn’t do you any good. Objective theories are attractive, in part, because of intuitions we have about the lives of people who are lacking certain key components of human flourishing (such as friendship). If these intuitions would vary depending on whether we are picturing our own life or someone else’s, which intuitions we choose to focus on could influence which theory looks most plausible.

¹According to Joshua Knobe (2005), “Experimental philosophy focuses on many of the same types of intuitions that have long been at the center of philosophical study, but it examines those intuitions using the methods associated with contemporary cognitive science – systematic experimentation and statistical analysis.”

1. The Pilot Study

The hypothesis I set out to test was that people have different intuitions about the contribution that knowing unpleasant truths makes to well-being depending on whether they are thinking of their own good or the good of someone they care about. In particular, I thought people would be more likely to want information for themselves and to think it makes them better off than they would be to choose information on behalf of a loved one or to think that the information would make that person better off. I gave subjects two scenarios and three different questions about each. My hypothesis was strongly confirmed in one scenario, but not in the other. Before I explain the experiment, I will describe the constraints I tried to follow in devising the scenarios.

Valuing Instrumentally vs. For its own sake

The cases had to be constructed so that the information has more than instrumental value. Obviously, being informed has instrumental value. If you know there's gasoline in the cup rather than gin you won't drink the liquid, and more of your desires will be satisfied in the future. But the claim that people value information or knowledge instrumentally as a means to meeting other goals is, to my mind, too obvious to be interesting. The more interesting question is about whether people value being informed for its own sake. This seems to be one of the intuitions behind Nozick's experience machine: there is something valuable about having real experiences, even if (by hypothesis) no more pleasure is gained in the real experience.

So, I think we have to ask people about the value of information in a way that moves their attention away from the instrumental value of being informed. One problem here is that there isn't an easy way to explain the distinction between two types of valuing. It won't work to say that the information "does you no good" or "has no benefit" because these phrases conflate intrinsic benefit and instrumental benefit. I want to know whether people think knowing the truth benefits them in a particular way, that is, as a constitutive part of their good. This is a difficult idea to explain. What I tried to do was to eliminate the influence of the most salient kinds of instrumental value: contributions to pleasure, felt happiness, or meeting particular goals.

Focus on Prudential Value

I wanted to ask about people's intuitions about prudential value or the non-moral good for a person. "Prudential value" is not a common notion in ordinary language, so I asked about what is "*good for*" a person and what makes a person "*better off*". I also asked about people's *preferences* because of the theoretical tradition that ties well-being to desires. There is some reason to think that the questions I have asked do not get at the concept of prudential value; I discuss this issue further in the final section of the paper.

The Scenarios

I gave 59 undergraduate students two scenarios, each followed by three questions and a space for explaining their answer or noting confusions. I've titled the two cases "Dour Dad" (DD) and "Madame Bovary" (MB). Thirty-one students received versions of the scenarios in which they were the subject of the case (the Self Condition). Twenty-eight students received versions of the

scenarios in which the subject was a friend or sibling (the Other Condition). Here are the cases, as they appeared:

Dour Dad (self)

Imagine that you are a husband and father who wants to be the kind of person who shares your feelings with your wife and children and can talk about problems in a productive, honest way. You really believe it's important to avoid being emotionally closed, and insensitive, as your own father was. And you have truly made your best effort. You think the efforts have paid off quite well; you think you are a warm and cuddly guy, very different from your father. In fact, though, you have made much less progress than you think. You never share your fears or concerns with your wife, and people (including your children) find you distant. Nevertheless, you're a good father and husband; your family loves you and they don't think this is a terrible problem. A good friend, someone whose opinion you always take seriously, sees that you haven't made as much progress as you think you have. If this friend were to talk to you about this, you would change your mind about how well you have met your goal.

Reflecting on this case, please circle your level of agreement with the following statements:

1. I would prefer that my friend tell me the truth, even though it would hurt.

○ ○ ○ ○ ○ ○ ○ ○
strongly disagree strongly agree

2. Even though I'm already doing my best, I would be better off in some way if my friend tells me the truth about how well my efforts have paid off.

○ ○ ○ ○ ○ ○ ○ ○
strongly disagree strongly agree

3. Just knowing the truth would be good for me, though it won't help me improve at meeting my goal.

○ ○ ○ ○ ○ ○ ○ ○
strongly disagree strongly agree

Madame Bovary (self)

Imagine that you are a widow or widower who was happily married for 10 years. Your spouse was attentive and loving; the two of you were emotionally close and shared just about everything. Having a good relationship was very important to you both and you were proud when friends commented that you had the perfect marriage. It turns out that there was one thing your spouse did not share with you: a love affair he or she had five years ago. The affair lasted a year, but you never had any evidence of it. Your spouse was very careful not to leave clues and your chances of ever finding out by accident are extremely small.

spouse was attentive and loving; the two of them were emotionally close and shared just about everything. Having a good relationship was very important to them both. You always thought that they had the perfect marriage and you know they were proud that people thought this. It turns out that there was one thing the spouse did not share with your friend: a love affair he or she had five years ago. The affair lasted a year, but your friend never had any evidence of it. The spouse was very careful not to leave clues and your friend's chances of ever finding out by accident are extremely small. Now that the spouse has died, the other person involved in the past affair comes to you for advice. He/she wants to know whether to tell your friend the truth. This person trusts you and will follow your advice without fail.

Reflecting on this case, please circle your level of agreement with the following statements:

1. I would prefer that the friend find out the truth, even though it would hurt.

strongly disagree strongly agree

2. My friend would be better off in some way if he/she knew about the affair.

strongly disagree strongly agree

3. Just knowing the truth would be good for my friend, even though it would make him/her feel unhappy.

strongly disagree strongly agree

Results

In the Madame Bovary case, there was a strong asymmetry between self and other.² (See Figure 1). On the question about preferences, the average response in the Self Condition was 4.23 and the average response in the Other Condition was 2.71, a difference of 1.52 on the 7 point scale. The other two questions also generated asymmetrical answers. The average responses for the “good for” question were 4.23 and 2.93 for self and other respectively. The results for the “better off” question were 3.65 and 2.71 for self and other respectively. Importantly for our purposes, the average responses for the preference and “good for” questions were above the neutral point (4) in the Self Condition (indicating a positive valuation of the information) and below the neutral point in the Other Condition (indicating a negative valuation of the information).

²The effect size for each question was moderate, ranging from one half (for the “better off” question) to two thirds of a standard deviation (for the other two questions). The results the preference and “good for” questions were statistically significant (p=.008 and p= .010, respectively). The results for the “good for” question were on the margin of significance (p=.072).




Figure 1

The results in the Dour Dad case were asymmetrical but not statistically significant. (See figure 1). Examining the students' comments about the Dour Dad case led to an explanation for the failure to confirm the asymmetry hypothesis in this case. Thirty-two of the subjects (15/31 in the Self Condition and 17/28 in the Other Condition) explicitly rejected the premise of this case, namely, that the information would not help the Dour Dad do better at meeting his goal. So, it seems clear that I failed to isolate intuitions about information valued for its own sake. This failure may explain the fact that the asymmetry between self and other was less dramatic. It likely also explains why the average responses (with Self and Other Conditions combined) were higher than in the Madame Bovary case.

I would like to be able to show that the self/other asymmetry has wider scope and is not limited to cases of marital infidelity. To do this, I need cases in which the instrumental value of information is better controlled (or, alternatively, measured). I will say more about this in section 3. For now, to make for the most interesting discussion, I will focus on the asymmetry that was found and the philosophical implications if it does turn out to have broader scope.

2. Philosophical Implications

The role of intuitions in normative theories is not straight-forward. In part, this is because when it comes to concepts like "prudential value" or "well-being", we are not quite trying to explicate our ordinary concept. We are trying to give an analysis – what we might call, following Dan Haybron (forthcoming), a "reconstructive analysis" – of a *normative* notion. In this context, one might argue that it is reflective, or considered judgments about cases (rather than unreflective intuitions) that are relevant to our analysis. Certainly this was Rawls's view in his discussion of

the method of reflective equilibrium, the most widely accepted method for defending normative theories (1971; see also Daniels 1979). Taking this point on board, it is nevertheless true that if intuitions about self and other pull in different directions, then they can't both be right. Another way of putting the point is that not all of the intuitions about the case of the cheating spouse can count as considered judgments. At least some of the subjects responding to the cases are making a mistake that should not be reflected in the development of a theory of prudential value.

The method of reflective equilibrium should lead us to think carefully about the conditions under which intuitions about prudential value are deserving of attention. With this method in the background, the self/other asymmetry in intuitions about prudential value prompts us to look for mistakes that explain this difference. Since the questions about the scenarios ask about the same things, the difference in answers must be explained by different assumptions about the subject of the questions (him/herself or the brother or friend). To explain the asymmetry, then, we should look for mistaken assumptions that skew our judgment in one or the other case. The thought here is that *considered* judgments would not be asymmetrical because proper reflection on the cases would eliminate the mistakes. If this is indeed the case, experiments that establish asymmetrical intuitions teach us that we have certain biases and that we need to be wary of trusting our intuitions insofar as this is the case.

Before we turn to a discussion of these biases, it is worth saying something more about the objection (discussed very briefly in the introduction) that intuitions are just entirely irrelevant to the project of normative theorizing. One could have this view for different reasons. First, one might believe that *considered* judgments are the place to start and so that we should just begin with philosophers' reflections and not worry about what "people think". In my view, this position overestimates the ability of professional philosophers to overcome the biases that afflict "people". Particularly if these biases have deep psychological explanations, philosophers may be just as likely to experience them and, therefore, to have asymmetrical intuitions. (In fact, in my conversations about these cases with other philosophers, I have heard very asymmetrical intuitions reported, even in the Dour Dad case. The reason for this anecdotal success in the Dour Dad case may be that philosophers are more willing to grant the assumption that the information has no instrumental value).

Second, one might have the view that it doesn't matter what people think even when they are being reflective, because one believes that normative theories have some other basis – say, a theory of human nature or practical reason. This is a more fundamental objection and one that I cannot address in much detail here. I will say that it has seemed to many that when it comes to prudential value or well-being we have more reason to care about intuitions than we do in the case of moral value. The thought here is that because prudential value is the good *for a person*, it must be related to the subject in an important way. We might say that it is a criterion of success for theories of prudential value in particular that they makes sense to us as capturing an intuitive notion of the good. In other words, the subjective character of prudential value may give us some reason to take intuitions seriously (even if we don't have such a reason in the moral case).

Moreover, as I said in the introduction, intuitions about cases seems to be at least partly responsible for which theories look attractive to us in the first place.

Finally, there may be practical implications of the fact that people make different judgments for their own case than they do for the case of someone they care about. These implications would be independent of the role of intuitions in normative theorizing.

Explaining the Asymmetry

In this section I discuss what I see as the two main contending explanations for the self/other asymmetry and comment on the implications of each by itself.

Self-Bias

One explanation for the self/other asymmetry is self-bias, or an inflated opinion of one's own capacity to process information appropriately. Given the large psychological literature on positive self-bias (e.g., Taylor and Brown 1994), it does seem likely that people could mistakenly believe that they will process the information more objectively than the average person, be less likely to be devastated by bad news, and so on.

In addition to the general self-bias literature, there is some experimental evidence in other domains that lends credence to this explanation. For example, in a study that investigated attitudes toward sources of information in consumer choices among brands of condoms, Wilson and Brekke (1994) found that although most students viewed information from Consumer Reports as more reliable than information from fellow students, they were more likely to want to see this information themselves than they were to want their peers to see it. Wilson and Brekke conclude that the students had "faith in their ability not to be influenced by it [student information] in an undesirable way" (1994: 125).

There is also some anecdotal evidence for this explanation from the student comments on the Madame Bovary case in the Pilot Study. In the Other Condition, several students mentioned the possibility that knowing about the affair would cause severe depression or exacerbate the grieving process as a reason not to tell the friend about his or her unfaithful spouse. Comments like these were not found in the Self Condition. This might reveal an overestimation of how well one is capable of handling difficult information.

Insofar as self-bias is the explanation for the self/other asymmetry, our intuitions about whether facing unpleasant truths is good for us when we reflect on our own case do not rise to the level of considered judgments. This point has implications for philosophical thinking about prudential value. First of all, unless philosophers are immune to (or able predictably to overcome) the effects of self-bias, which seems unlikely, the self/other asymmetry gives us a reason to be skeptical about our self-regarding judgments about prudential value. Second, the self-bias explanation indicates a particular way in which our desires for our own case are not good indicators of what's good for us, even by our own lights. From a practical point of view, this could lead us to the view that when we're wondering whether facing some painful truth would be

good for us, we ought to try to think of ourselves from the third person point of view, as we would think of a dependent or ward. From a theoretical point of view, the self-bias explanation underscores a problem for desire theories of prudential value.

Attributing Values

A second explanation for the asymmetry has to do with the assumptions we make about others in order to decide whether something is good for them. If when we make such decisions we make assumptions about what a person wants or cares about, and if we are inclined to be cautious in the assumptions we make with respect to others' good, then an asymmetry could result, because the assumption that a person wants or cares about the truth despite the fact that it will only bring pain seems like a risky assumption.

This explanation relies on several assumptions, but they are plausible ones. First, this explanation assumes that we make assumptions about others' desires or values when we think about what's good for them. Given the prevalence of desire and preference as a key component of theories of well-being, this seems like a plausible assumption. Second, it assumes that we are cautious with respect to the assumptions we make about others' good. Some evidence for this comes from the literature on risk imposition, according to which we are less likely to impose risks on others than we are to take them on ourselves (Fernandez-Duque and Wifall 2007). In the present case, the thought is that by imposing painful truth on someone we risk their happiness for the chance that they would prefer to know despite the pain. In our own case, we believe we know how much value we assign to knowing the painful truth so there is little perceived risk in opting to have the information (when we claim to value it). In the case of another person, however, I think risk aversion would weigh against attributing an attitude of intrinsic valuation of painful truth to another person. Thus the third assumption – namely, that people perceive it to be less risky to attribute the value of happy feelings than the value of painful truth to others – also seems plausible.

If this is the right explanation for the self/other asymmetry, the decision about which intuitions rise to the level of considered judgments seems to depend on whether our assumptions about what others want or care about are correct. We can get some information relevant to this question from the Pilot Study. According to the Pilot Study, people do differ in the degree to which they prefer painful truth even in their own case: the distribution of responses was bimodal. (See figure 2). Interestingly, responses in the Other Condition were unimodal and negatively skewed. (See figure 3). It seems that people are unlikely to attribute an interest in knowing painful (but non-instrumentally valuable) truths to others. According to the self-reports as shown in figure 2, however, some people do have such an interest. If this is right, then the explanation of the self/other asymmetry is a kind of misinformation about the desires or concerns of other people, or a reluctance to attribute certain kinds of values to others. What are the philosophical implications of this explanation?

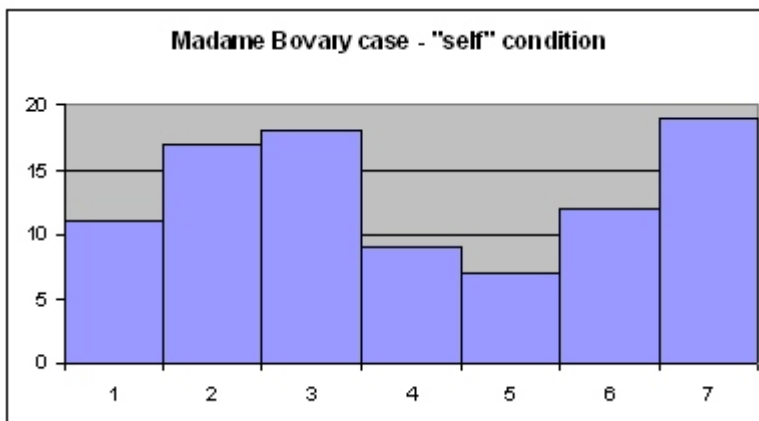


Figure 2: this graph takes all three questions averaged together. Numbers on the vertical access represent number of respondents.

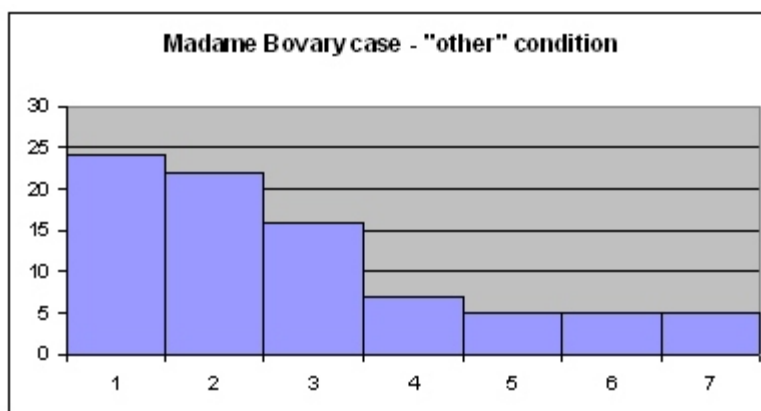


Figure 3: as in Figure 2, the three questions are averaged.

One thing we might conclude is that our other-directed intuitions are not necessarily a good guide to what is prudentially good. If we do tend to base our intuitions about what is good for another person on assumptions about that person's desires and concerns, and if we also tend to make false assumptions about those desires and concerns, our judgments about what is good for them do not count as considered judgments.

As was true with respect to the self-bias explanation, this point has both practical and theoretical implications. From a practical point of view, if we think that what others want or care about is important to what's good for them, when we make decisions on behalf of others we would do well to think carefully about the assumptions we make about such matters. From a theoretical point of view, the fact that we make cautious attributions in this context might have implications for the defense of objective theories of the human good that ask us to rely more heavily on our intuitions about people we care about. Further, this explanation might have implications for the correct characterization of an ideal advisor. Typically, in desire-based theories of well-being that

follow an advice model, the ideal advisor is characterized as fully informed. Defining what counts as “full information” has proved to be difficult, but one thing we can learn from the present study is that a fully informed advisor would have to overcome the human tendency not to attribute certain values to those she is advising.³

Conclusion: Putting the two explanations together

I have discussed two possible explanations for the self/other asymmetry with respect to learning painful truths. Each explanation reveals a different problematic assumption on the part of the subject answering questions about the scenarios, one about the self and one about other people. If one of these explanations were correct to the exclusion of the other, we would have some reason to discount one kind of intuition and weigh the other one more heavily. But these explanations are not mutually exclusive. The same person could overestimate her own ability to handle painful truths and make incorrect attributions of desires for information on behalf of other people. Given the evidence, it seems plausible that both explanations are true to some extent. What this means for purposes of defending claims about prudential value is that we have reasons to be skeptical about intuitions whether they are about oneself or another. On this line of thought, both our intuitions about ourselves and about others are problematic; neither is a perfect guide to prudential value.

The value of an experimental investigation into our intuitions, then, is that it can reveal where we have systematic biases, which can help inform us about what we need to correct for when we rely on our intuitions. When we ask the question “Is it good for us to face the painful truth about ourselves and our relationships?” we would do well to reflect on our judgments about our own lives and the lives of those we care about and to be mindful of the biases we bring to both kinds of judgment. This is true whether we are defending theoretical claims about prudential value and well-being, thinking about out how to live own lives, or figuring out the best practical advice for others.

Finally, for those of us who might be inclined to investigate the question about painful truths empirically, there is a related lesson to be learned. We can ask people for their views about what’s good for them in different ways, each of which may be subject to its own distortion.⁴ If

³It seems clear that a *fully informed* advisor, as this has been understood by, e.g., Railton (1986), would know the actual concerns of the person she is advising. I mention this point because it may be of interest to the development of a more realistic notion of an ideal advisor, one that does not result in total alienation between subject and advisor.

⁴Matt Adler and Paul Dolan are working on a survey project that raises this concern: “We’re asking people to rank possible lives --- specified in terms of income, health, happiness, and longevity -- in terms of their self-interested preferences. (“Which life would you prefer to lead?”) The hope is that surveys, along with philosophical reflection, will help us figure out what the dimensions of well-being (ie., objective goods, or aspects of worlds that people generally prefer) are, and eventually to calibrate a cardinal utility function”. (Matt Adler, email

the self/other asymmetry does have broad scope (which remains to be shown), the best strategy might be to ask different questions that require people to think about the prudential good from different perspectives. Doing so may lead us to the truth by helping us to identify the reasons for the variation in responses and to eliminate all forms of undesirable bias.

3. Future Work

I plan to run the experiment again, using only the Madame Bovary case. I would also like to experiment with different cases that will allow me to infer that the asymmetry has greater scope. I would greatly appreciate feedback about three problems I anticipate as I proceed.

First, I'd like to make sure the questions I ask the subjects are capturing the notion of *prudential* value. It was interesting to me that in both cases (DD and MB) the "better off" question turned up the smallest asymmetry. I'm not sure why this is. Is "better off" connoting something economic? I also worry that "good for" might connote something moral. What about asking subjects to think about what they would do "for someone's *sake*"? Or, I could ask a question about what kind of *life* they would prefer to lead. (I think this would require significant changes to the scenario, though).

Second, I'd like to know what other cases will have the best chance of producing the self/other asymmetry. One case I thought of that I think would be compelling is of a parent of a soldier who has to think about whether she would want to know that her son or daughter was killed by friendly fire rather than in battle (or, perhaps better, whether her son or daughter died painfully). (I've included a few other cases I had included in my first presentation of this work in Appendix A).

What I would really like, though, is a case in which it is self-knowledge that is at issue. As evidenced by the difficulty I had with the Dour Dad case, I have found this hard to do, because the instrumental value of self-knowledge is so obvious. Another example of a case that involves self-knowledge is "You're not as great as you think you are" (see Appendix A). Again, it's very difficult to control for instrumental value here, but some variation on this case might work. One solution to this problem would be to *measure* the degree to which subjects are relying on the instrumental value of self-knowledge (since it seems to difficult to control for it). I don't know whether this would be a good idea.

Third, I'm wondering if it would be a good idea (eventually) to add questions that test the two explanations I have proposed. For example, I could ask people questions that would assess the degree of self-bias, or their assumptions about other people's desires and concerns.

correspondence).

Appendix A

You're not as great as you think you are (self)

Psychologists have shown that people exaggerate their own talents and abilities so that just about everyone thinks that he or she is a better driver, student, friend, and so on, than he or she actually is. Given this research, it's quite likely that you have an inflated opinion about yourself in at least some areas of life. Assume for a minute that this is true: you do think you're a better student than you really are. You're not as talented at writing as you think you are, and your interpretations of books and films are not as original and creative as you think they are. You can't really know what the consequences would be of knowing the truth – the results could be good or bad. Putting the consequences aside, then...

Deluded achiever (self)

You have a wonderful memory of standing up to the schoolyard bully when you were 6 years old. The bully was picking on your friend, a much smaller kid, and you proudly remember yelling at him to pick on someone his own size and threatening to whack him with your lunchbox. The bully backed off, you recall, and left you and your friend alone for the rest of the school year. As it happens, this isn't quite how it happened. In fact, you were just a bystander to the event and it was really your friend who stood up to the bully. The change in memory first started when you reported the story to your older sister, exaggerating your own role to impress her. Over the years, your memory has gotten embellished and further distorted. The false memory is a nice one and it hasn't caused any damage. The kid who really stood up to the bully moved far away and the two of you have no mutual friends. One of your good friends was there at the scene and knows the truth.

Gossipy friends (self)

Two of your good friends sometimes talk about you behind your back. They do it because they're a little bit jealous of you and this makes them feel better. Mary and Jane really are good friends; the gossip is a bit mean, but it won't come to any harm. They only gossip with each other (they don't spread rumors about you to other people) and their jealousy doesn't prevent them from being considerate friends. Putting aside what you would do with the information,

Works Cited

- Daniels, N. 1979. "Wide Reflective Equilibrium and Theory Acceptance in Ethics" *Journal of Philosophy* 76:5:256-82
- Feldman, Fred. 2004. *Pleasure and the Good Life: Concerning the Nature, Varieties, and Plausibility of Hedonism*. (Oxford: Oxford University Press).
- Fernandez-Duque, D. and Wifall, T. 2007. "Actor/observer asymmetry in risky decision making", *Judgment and Decision Making*, vol. 2, no. 1, pp. 1-8.
- Haybron, D. Forthcoming. *The Pursuit of UnHappiness*, forthcoming, *Oxford University Press*.
- Knobe, J. 2007. "Experimental Philosophy", *Philosophy Compass* 2 (1), 81–92.
- Nozick, R. 1974. *Anarchy State, and Utopia* (New York: Basic Books).
- Railton, P. 1986. "Moral Realism". *The Philosophical Review* XCV.2 (April):163-207.
- Rawls, J. 1971. *A Theory of Justice*, 2nd Edition 1999, Cambridge, MA: Harvard University Press
- Sumner, L. W. 1996. *Welfare, Happiness and Ethics* (Oxford: Clarendon Press, 1996).
- Taylor, S. E. and Brown, J. D. 1994. "Positive Illusions and Well-Being Revisited Separating Fact From Fiction". *Psychological Bulletin* 116.1 (July): 21-27.
- Wilson, T. D., & Brekke, N. C. 1994. "Mental contamination and mental correction: Unwanted influences on judgments and evaluations". *Psychological Bulletin*, 116, 117-142.