
How to Think About Satisficing

Chris Tucker

College of William and Mary

(Forthcoming in *Philosophical Studies*)

1. Introduction

Suzy doesn't really care about getting as much of the good as she can. She sometimes rejects the better for the good enough because the (mere) good enough gets her what she really wants. We can say that Suzy is a *satisficer*. Mary, in contrast, *does* care about getting as much of the good as she can. Nonetheless she sometimes rejects the better for the less good because she has a countervailing consideration, i.e., getting the better would require her to give up something else she cares about, something that is, in some sense, independent of how good her options are. Perhaps, for example, the better choice violates someone's rights. We can say that Mary is a *motivated submaximizer*.

Satisficing and motivated submaximization, so characterized, are different ways of choosing a suboptimal option. Yet this difference is easily missed. In fact, defenders and detractors of satisficing may, to some extent, be talking past one another. Putative arguments for appropriate "satisficing," if they are successful at all, tend to support only that motivated submaximization can be appropriate.¹ In contrast, objections to satisficing tend to be aimed, not at motivated submaximization, but at what I call (*radical*) *satisficing*.

The purpose of this paper is two-fold: (i) to clarify the distinction between (appropriate) motivated submaximization and (appropriate) satisficing and (ii) to show that the three most prominent objections to satisficing fail. The result will be a kind of stalemate: I show that currently there are no good arguments *for or against* the claim that satisficing can be appropriate. I'm sympathetic with satisficing, but a positive argument that satisficing can be appropriate must wait for another day.

The first two sections of the paper clarify the distinction between satisficing and motivated submaximization. I explicate motivated submaximization in section 2 and satisficing in section 3. The labels here do not matter. What matters is that we have some vocabulary to keep track of the two different ways of choosing a suboptimal option. If you prefer to think of motivated submaximization as a type of satisficing, then let my use of *satisficing* be shorthand for *radical satisficing*.²

The remaining sections of the paper consider arguments for and against the claim that satisficing can be appropriate. In section 4, I explain why the existing justifications for choosing a suboptimal option justify motivated submaximization, but not satisficing. In sections 5, 6, and 7, I consider and disarm the three most prominent objections to satisficing. Sections 5 and 6 are especially important, because they construct a vocabulary that allows us to understand and track the differences between appropriate motivated submaximization and appropriate satisficing. This discussion will lead us to

¹ I tend to use the term *appropriate* and its cognates (rather than *rational* or *morally permissible* and their cognates), because nothing in this paper hinges on the difference between moral and rational evaluation.

² With that said, in note 17, I explain why I don't think the label *satisficing* should be extended to motivated submaximization.

refine our understanding of *countervailing consideration*, but until then, you can rely on your intuitive understanding of the term.

2. Motivated Submaximization

Satisficing and motivated submaximization are kinds of submaximization. An agent **submaximizes** iff she chooses an option when she had a better option. In other words, you submaximize when you choose a **suboptimal option**, an option not as good as one of your other options. Ignorance can lead to appropriate submaximization. When you don't know what all of your options are, the best option that you are aware of may not be your best option overall. In such a case, the consensus is that you can appropriately choose the best option that you are aware of, provided that the option is "good enough" (Petit 1984: 166–7; Schmidtz 1995: 29–30; Byron 1998: 71–4; Richardson 2004: 106). Perhaps submaximization can also be appropriate when you reasonably but mistakenly think you are choosing the best option. I set aside such epistemic excuses by focusing on a special kind of choice situation.

It is convenient to assume that, at any given time, an agent has a determinate list of options. The individuation of options is controversial business, but the rough idea is that the agent's options capture states of affairs that can be brought about by the agent in those circumstances.³ Sometimes one of your options may be to search for more options. If you are trying to decide whether to accept an offer on your home, one of your options may be to wait and see if other offers come in. But if you do have the option to wait, that option will be evaluated along with all other options. It might be your best option, your worst, or something in between.

A's situation is **transparent** iff (i) A knows precisely what A's options are, and (ii) A knows the ranking of all A's options. Transparent situations ensure that, if the agent chooses an option that is suboptimal, then she does so with full knowledge. In transparent situations, a choice for less than the best is never made in ignorance. Epistemic excuses don't apply.

We can get a handle on what I call *motivated submaximization* by considering the most plausible and widely held explanations for why submaximization (in a transparent situation) can be appropriate. It is relatively uncontroversial that submaximization is appropriate whenever there is some countervailing consideration that motivates it. I will briefly discuss four of the most familiar motivations for submaximization.

The first motivation concerns Ever Better Situations. In these cases, one has no optimal option because, for every option one can choose, there is another better. Suppose that a genie offers to ensure that your life enjoys any degree of well-being you choose. Since there is no maximum degree of well-being that you can enjoy, for every degree of well-being you can have, you could have some degree higher. Sorenson (2006: 214) contends that the lack of an optimal option guarantees that your choice is irrational. While there is something to be said for Sorenson's position, it has seemed intuitively obvious to many philosophers that, in Ever Better Situations, you can rationally choose some arbitrarily high degree of well-being for yourself (e.g., Pollock 1983: 417-8; Schmidtz 1995: 42-45, 2004: 41-4; Slote 1989: 110-23; and Langtry 2008: 74-8).

³ Those in the satisficing literature often take controversial stands as to what counts as an option. For some examples, see Pettit and Brennan (1986: 440-1); Slote (1989: 24-5), and Henden (2007: 340, nt 1).

A second motivation for submaximization concerns your special connection or relationship with specific individuals (Cottingham 1986; Hurka and Shubert 2012; Scheffler 2010; Vallentyne 2006: 25; and van Roojen 2004: 170). Suppose you have two options, A and B. A best promotes the well-being of your family while also making everyone else better off than they otherwise would have been. B is good, though not as good, for your family; it does, however, best promote the well-being of everyone else. One might agree that B is the best option overall since it provides the greatest benefit to the greatest number but then deny that one must choose it: your special connection to your family makes it appropriate for you to choose A even though you recognize it is less than the best.⁴ This motivation, like the first, is compatible with aiming at as much good as one can get. The person who appropriately chooses suboptimal A might aim at doing the greatest good for the greatest number. Yet, when this aim conflicts with her aim of doing the greatest good for her family, she can appropriately choose less than the best. Even critics of this position recognize that it is plausible enough to be part of “ordinary, commonsense morality” (e.g., Kagan 1989: xi, 2–3).

A third motivation for submaximization concerns incommensurability. Suppose goods G1 and G2 are incommensurable. Option 1 best promotes G1 and is very good with respect to G2. Option 2 is very good with respect to G1 and best promotes G2. Since G1 and G2 are incommensurable, there is no broader perspective from which we can say that one of the options is overall better than the other (or even that they are equally good). Suppose I choose Option 1. I thereby choose an option that is suboptimal with respect to G2. I might make this choice even though I aim at getting as much of G2 as I can; I just accept a suboptimal amount of G2 because I have a competing aim, namely that I’m also aiming to get as much of G1 as I can. Those who appeal to incommensurability to justify submaximization include Schmidtz (1995: 45-50; 2004: 44-8), Richardson (2004), and Weber (2004).

A fourth motivation concerns deontological side constraints on the promotion of goodness (Vallentyne 2006: 28-33; van Roojen 2004: 170). For example, I might aim at doing the greatest good for the greatest number but nonetheless choose a suboptimal option because it is the best option that doesn’t violate anyone’s rights.

We’ve identified, in this section, four widely held motivations for submaximization. If we pay careful attention to the four motivations, we can see that each of them has a common structure. Each motivation is compatible with the agent’s aiming at as much of the good as she can get. What made choosing a suboptimal option plausibly appropriate in the above cases is that there were independent considerations that conflicted with that aim; there were countervailing considerations that (allegedly) made choosing a suboptimal option appropriate. Let’s use the following term to capture this shared feature:

⁴ *Objection:* the special connection with your family should be factored in when we evaluate options. Once it is factored in, A is better than B. *Reply:* the objection presupposes a particular axiology that isn’t always endorsed by those who think special connections matter (e.g., Vallentyne 2006: 30-2). Furthermore, we should avoid assuming that everything that makes a difference to the quality of a choice must also make a difference to the quality of our options (van Roojen 2004: 163-9, especially 163, 166-7).

Motivated Submaximization: an agent A *submaximizes with motivation* in a transparent situation iff

- (a) A aims at getting as much of good G as A can, but
- (b) A chooses a suboptimal option with respect to G because of some countervailing consideration.

Each of the above motivations for choosing a suboptimal option counts a different kind of motivated submaximization that is widely thought to be appropriate. It's popular, then, to hold that motivated submaximization can be appropriate.

3. Satisficing

Satisficing, as I define it, is a type of *unmotivated* submaximization (cf. Pettit 1984: 174). One **submaximizes without motivation** iff she chooses a suboptimal option but *not* because she is motivated by some countervailing consideration. Don't read too much into the distinction between motivated and *unmotivated* submaximization. The distinction is shorthand for, roughly, 'submaximization motivated by countervailing considerations' and 'submaximization *not* motivated by countervailing considerations', respectively. When I say that some choice is *unmotivated*, I'm not denying that it is motivated by reasons; I'm denying that it's motivated by reasons of a special sort, namely reasons that countervail other reasons. It will become clear, especially in section 6, that appropriate satisficing always involves reasons of a certain kind.

To *satisfice* is to act with a certain aim, which distinguishes satisficing from both motivated submaximization and unmotivated submaximization more generally. One doesn't *satisfice* insofar as one aims at the best or getting as much of the good as one can. One doesn't *satisfice* if she chooses the good enough for no reason whatsoever. To *satisfice* is to choose a suboptimal option because one is aiming to promote some good to degree D but not as much as one can.

To **promote** some good to degree D is to take the necessary means of bringing about a degree at least as high as D. If an agent aims at promoting G1 to degree D when D is something like *as much as I can get* or is *the most possible*, let us say that the agent **aims at the optimum**. If an agent aims at promoting G1 to degree D when D is something short of the optimum, let us say that the agent **aims at the good enough**. Optimizers aim at the optimum and satisficers aim at the good enough. Their aims are common insofar as they aim at the promotion of value. Their aims are distinct insofar as they aim at differing degrees of value.

Satisficing, then, involves choosing a suboptimal option because it realizes one's aim at the good enough. Yet if we want to capture an interesting and controversial notion of satisficing, we must impose further conditions on the aim at the good enough. Suppose I'm aiming only at a good enough degree of momentary well-being because I know doing so will maximize my overall well-being. Even self-proclaimed critics of satisficing allow choosing the suboptimal option as a *means* to optimizing some other good to be appropriate (e.g., Byron 1998: 80-1; Richardson 2004). So our conception of satisficing must add that the aim is taken for its own sake.

Even with this constraint we have not captured a type of choice that is particularly controversial. Some ends are more final than others.⁵ As Pettit and Brennan (1986)

⁵ Richardson (2004: 119-23 and nt 40 on pg 129) provides some helpful discussion concerning how ends, or aims, can be arranged in a hierarchy, with some ends being more final than others.

argue, I might aim at the good enough for its own sake as a way of aiming at the optimum for its own sake. My ultimate aim in life might be to get as much of the good as I can. I then might learn that I can achieve this aim only if I develop the habit of aiming at the good enough for its own sake. I might develop such a habit. I'm not satisficing if I then choose a suboptimal option because it realizes my aim at the good enough. Satisficing requires that one aim at the good enough *purely* for its own sake, that it not be a means to any further aim. In other words, one's end of the good enough must be an ultimate, or purely final, end (cf. Slote 2004: 17-18).⁶

I now present the official definition of satisficing:

Satisficing (with respect to good G1): an agent *A* *satisfices* with respect to G1 in a transparent situation iff:

- (a) *A* aims, purely for its own sake, at promoting G1 to degree *D* but not as much as *A* can, and
- (b) *A* chooses a suboptimal option with respect to G1 that has a value greater than or equal to *D* because *A* knows it satisfies the aim in (a).

This notion of satisficing is relativized to good G1. If there is more than one (intrinsic) good—for example, if both well-being and beauty are intrinsic goods—the relativization allows that one satisfices with respect to good G1 without satisficing with respect to good G2. A transparent situation, recall, is one in which an agent knows what her options are and how to rank them. The agent who satisfices in a transparent situation *knowingly* rejects the better for the good enough.

To accuse someone of satisficing, then, is to accuse them not only of rejecting the better for the good enough, but also of having a certain motivational structure. The satisficer aims at promoting the good to some suboptimal degree purely for its own sake. She chooses the good enough because it realizes this aim. The motivated submaximizer has a different motivational structure. She aims at promoting the good as much as she can. She chooses a suboptimal option, because she has some special consideration or competing aim that motivates the choice to reject the better for the good enough.

Some of those who claim to defend the appropriateness of satisficing in some interesting sense do not defend the appropriateness of this sort of satisficing (e.g., Dreier 2004; Greenspan 2009; Narveson 2004; van Roojen 2004: 170-1; and Weber 2004: 98). In such cases, these theorists usually defend the appropriateness of only some sort of motivated submaximization. Even classic defenses of satisficing, such as Slote's 1989 and Swanton's 1993, fail to count as defenses of satisficing in my sense. They defend something much closer to motivated submaximization, and Slote (1989) seems to endorse the first objection to satisficing discussed below (see sec 5).⁷ Putative

⁶ Those who deny that satisficing can be appropriate tend to use the language of *global* ends as opposed to *purely* final ends (see, e.g., Schmitz 1995: 45-6; 2004: 44-5; Byron 1998: 76-80). I resist using the term 'global', because I tend to think of a global end as one that encompasses all my other ends. Schmitz (1995: 46-50, 2004: 45-8) stresses that, as he uses the term *global*, one might have multiple global ends which compete with one another.

⁷ This corrects my earlier "considered opinion" that Slote and Swanton did defend satisficing (2016: 134, nt. 7).

defenses of satisficing are almost always defenses of motivated submaximization (or something in that neighborhood).⁸

4. The Greater Plausibility of Motivated Submaximization

Now that we understand the difference between satisficing and motivated submaximization, we can evaluate these types of choice. There are several widely endorsed justifications for motivated submaximization, and so it's plausible that motivated submaximization can be appropriate. On the other hand, there are no existing arguments that justify satisficing. The arguments that have been furnished on satisficing's behalf are better construed as arguments for the appropriateness of motivated submaximization. Or, at best, these arguments show only that *either* satisficing *or* motivated submaximization can be appropriate.

Recall from section 2 that motivated submaximization is supported by the four most popular justifications for knowingly choosing a suboptimal option, namely the absence of an optimal option, special connections to specific individuals, incommensurability, and deontological side constraints. In contrast, there are no existing arguments that provide support for the appropriateness of satisficing. Consider, first, why the four considerations that allegedly justify motivated submaximization cannot justify satisficing.

Satisficing involves aiming at the good enough, and it's doubtful that the relevant considerations provide a reason to have such an aim. For example, a reason to prioritize my family's well-being over global well-being is *not* a reason to aim at the good enough with respect to global well-being. One aim can be prioritized over another even if both aims concern the optimum for some good. Furthermore, even if you think that these considerations do provide some reason to aim at the good enough, they do not provide a reason to aim at the good enough *purely for its own sake*. If a side constraint against torture gives me reason to aim at the good enough with respect to overall well-being, it does not give me reason to aim at the good enough purely for its own sake. Rather, it is for the sake of the side constraint that I ought to aim at the good enough. More generally, since each one of these four considerations is distinct from the good to be promoted, they apparently can't serve as reasons to aim purely for its own sake at any particular degree of the good, whether it be at the good enough or the optimum.

Supererogation is regularly discussed in connection with satisficing (e.g., Rogers 2010; Slote 1989: 29; Vallentyne 2009: 27-8). If supererogation is possible, then it is possible to appropriately and knowingly reject a better option for one that is good enough. It does not follow, however, that one can appropriately and knowingly reject a better option for one that is good enough *even if* one lacks a countervailing consideration. The mere possibility of supererogation does not show us that it might be appropriate to reject the better for the good enough because one aims purely for its own sake at the good enough. At most, the possibility of supererogation would show us that either satisficing or motivated submaximization can be appropriate. It does not provide any reason to think that satisficing, in particular, can be appropriate.

Slote discusses cases in which it seems that one can appropriately choose a suboptimal option (e.g., appropriately rejecting a beneficial snack, appropriately accepting a less than maximal offer on one's home) and these cases are sometimes

⁸ Vallentyne (2006: 21, 27-8) and Rogers (2010) may be exceptions to this rule.

thought to provide support for satisficing. In their current form, however, they support the idea that either satisficing or motivated submaximization is appropriate.⁹ Indeed, Slote himself interprets these examples as justifying something much closer to motivated submaximization than satisficing.¹⁰

Motivated submaximization has a lot going for it and, so far, satisficing has nothing going for it. Yet I write as satisficing's advocate, not its accuser. My goal in the rest of the paper is to provide a limited defense of satisficing. I will clarify the commitments of appropriate satisficing and, thereby, show that the three most prominent objections to appropriate satisficing are failures.

5. The Insufficient Reason Objection

5.1. The Objection

The Insufficient Reason Objection may be the most widely endorsed objection to satisficing. Even Slote may endorse it. Its first premise is:

Sufficiency: Whenever option O1 is better than a “good enough” alternative O2, we have some reason to choose O1 over O2.

Slote (1989: 39) endorses this premise when he holds that we have a “standing reason” to choose the best. More generally, Sufficiency is very popular especially among critics of satisficing.¹¹

The second premise is:

Necessity: When one has a reason to choose O1 over O2, one's choice for O2 is appropriate only if one has a countervailing consideration.

Slote (1989: 22) seems to have Necessity in mind when he maintains that the moderate agent chooses the best, other things being equal. Even if Slote doesn't endorse Necessity, it too is very popular among critics of satisficing.¹²

The conjunction of these two premises entails that satisficing, in the sense explicated in this paper, is necessarily inappropriate. According to Sufficiency, having better options is sufficient to have a reason to choose those options over the less good ones. When you have reason to choose the better options, Necessity says that you need a countervailing consideration to appropriately choose a less good option. The

⁹ Although I won't argue the point here, I think this argumentative strategy can be modified to support the appropriateness of satisficing.

¹⁰ Slote's moderate agent has a standing reason to choose and aim for the highest degree of well-being that she can get (1989: 22). So, in the absence of countervailing considerations, the moderate agent chooses and aims for the best (22). But the moderate agent has a countervailing consideration, namely “non-need,” and thus may very well not aim for the best and may very well choose less than the best (1989: 39; cf. Swanton 1993: 40-1). Non-need countervails well-being, because its reason-giving force operates against and independently of the reason-giving force of well-being (1989:36; cf. 32-33 and Swanton 1993: 37-8). Since Slote holds that choosing and aiming for less than the best requires countervailing considerations, he holds that satisficing cannot be appropriate.

¹¹ See, e.g., Byron (1998: 85); Kagan (1989: 61); Henden (2007: 342); Pettit (1984: 172); and Schmidtz (1995: 39, 2004: 39).

¹² Actually, what's popular among critics is even stronger than Necessity. Critics tend to claim that, if O1 is (all things considered) better than O2, then it's *impossible* (not even countervailing considerations can help) to appropriately choose O2. Kagan (1989: 379-81) endorses Necessity rather explicitly, but even he endorses the stronger claim too. Others who endorse the stronger claim include Henden (2007: 348) and Schmidtz (1995: 38-9; 2004: 38-9). Compare with Byron (1998: 79).

conjunction of Sufficiency and Necessity entails, therefore, that choosing a suboptimal option is appropriate only when one has a countervailing consideration. Satisficing, given this conjunction, is necessarily inappropriate. For when an agent satisfices, she chooses a suboptimal option because O2 realizes her aim at the good enough, not because she has a countervailing consideration.

5.2. *My Reply*

I endorse Sufficiency, the claim that if O1 is better than O2, then we have some reason to choose O1 over O2. The objection fails, because we can reasonably resist Necessity, the claim that when one has a reason to choose O1 over O2, it is appropriate to choose O2 only if one has a countervailing consideration. To understand why, we need to consider the distinction between requiring and justifying strength.

A reason has **requiring strength** with respect to choosing option O iff it makes *not* choosing O pro tanto *inappropriate*. A choice is **pro tanto inappropriate** iff it is inappropriate in the absence of (sufficiently strong) countervailing considerations. The fact that some surgery would cause me great pain is a reason that pro tanto requires me *not* to have the surgery. In other words, the great pain is a reason that makes it pro tanto *inappropriate* to have the surgery. This pro tanto requirement fails to be an all things considered requirement iff there are relevant countervailing considerations, e.g., the surgery is the only way to save my life. In the absence of countervailing considerations, it is inappropriate to ignore reasons with requiring strength.

A reason has **justifying strength** with respect to choosing some option O iff it has the power to make choosing O pro tanto appropriate. A choice is **pro tanto appropriate** iff it is appropriate in the absence of countervailing considerations with (sufficiently strong) requiring strength. The benefit of eating the snack has justifying strength, it is a reason that makes it pro tanto appropriate to eat the snack. This pro tanto appropriateness fails to be all things considered appropriateness iff I have relevant countervailing considerations with requiring strength, e.g., if eating the snack would leave my hungry child with nothing to eat.¹³

I assume, though nothing hinges on this, that to be a reason is to have *justifying* strength. In other words, R is a reason to choose O iff R makes it pro tanto appropriate to choose O.¹⁴ It does not follow that all reasons have *requiring* strength. That is, it does not follow that all reasons to choose O make it pro tanto *inappropriate* to *not* choose O. A **merely justifying reason** for choosing O is a reason that makes it pro tanto appropriate to choose O but does not make it pro tanto *inappropriate* to *not* choose O. Merely justifying reasons pro tanto justify without pro tanto requiring a choice. As such, they can be appropriately ignored even in the absence of countervailing considerations. That's just what it is to be a reason that pro tanto justifies but does not pro tanto require. The existence of such reasons is, as you might expect, controversial.

The proponents of the Insufficient Reason Objection deny that such reasons exist. This shouldn't come as a surprise, since this denial follows from Necessity. If a reason

¹³ My distinction between justifying and requiring strength is influenced by that of Gert (2007a, 2016), but I intentionally deviate from his characterizations.

¹⁴ To be sure, I allow that something can pro tanto require without pro tanto justifying. I just prefer to think of such things as providing coherence constraints on which choices are appropriate rather than as providing a reason to make a choice. But nothing of consequence hinges on this preference.

can't pro tanto justify choosing O1 unless it pro tanto requires it, then countervailing considerations are *always* needed to appropriately choose an alternative option O2.

The proponents of the Insufficient Reason Objection treat the impossibility of merely justifying reasons as some obvious axiom of rationality that needs no defense. In the satisficing literature, the most impressive defense of this denial is (not joking) Kagan's (1989: 380-1) remarks that it is "hard to imagine" and "hard to understand" how a reason could be merely justifying.¹⁵ At least he recognizes that there is an issue here. Yet merely justifying reasons are neither hard to understand nor hard to imagine. What we have is two distinct roles: a role of making something pro tanto appropriate and a role of making something pro tanto inappropriate. It is easy to understand what it means for one thing to play the former role without playing the latter. And, as the forthcoming discussion shows, it is easy to imagine something playing the former role without playing the latter one. Gert's work is also relevant here. He (2007a, 2016) vigorously defends the idea that justifying strength can outstrip requiring strength, and he also occasionally argues for the stronger claim that there are merely justifying reasons, reasons that justify without having any requiring strength at all (Gert 2000).

What the proponent of the Insufficient Reason Objection needs—and does not have—is an argument that nothing can play a pro tanto justifying role without its also playing a pro tanto requiring role. Or at the very least, he needs some reason to deny that the two roles come apart in the way that is required to vindicate satisficing.¹⁶ Since the objector doesn't have such an argument, they provide no reason to believe their key premise, Necessity, and their objection fails. My goal in the rest of this section is to articulate what I take to be the central theoretical claim of those who hold that satisficing can be appropriate. If what I call *Asymmetry* is this central theoretical claim, then the Insufficient Reason Objection has an even more serious problem.

Assume that merely justifying reasons exist. Gert (2007a: 535) argues that the existence of such reasons would disprove satisficing theories of rationality. The actual connection is nearly the opposite of what Gert suggests: the existence of merely justifying reasons is necessary—but not sufficient—for appropriate satisficing. Consider, for example, a view that grants that we each have a merely justifying reason to pursue our own projects, regardless of whether these pursuits would bring about the most good (cf. Kamm 1996: 230; Scheffler 1982, ch 3). These merely justifying reasons would countervail the reasons provided by overall goodness. Such a view allows a kind of *motivated* submaximization to be appropriate without allowing satisficing to be appropriate.

¹⁵ Kagan's term for a merely justifying reason is *noninsistent reason*.

¹⁶ Two objections to the possibility of merely justifying reasons are worth mentioning. First, if some sort of optimizing conception of appropriateness is true, then I think merely justifying reasons are impossible. Yet if we want an objection to merely justifying reasons to also pose an interesting objection to the possibility of appropriate satisficing, we can't start by assuming optimizing theories of appropriateness. Satisficing theories of rationality/morality are being put forward as competitors to optimizing theories. Second, Tenenbaum (2007: 162-4) raises an interesting objection to Gert's discussion of merely justifying reasons. I won't discuss the objection here, except to remind Tenenbaum that there are coherence constraints on rationality (see my nt 14 above). Even if one could permissibly have been in either of two states (e.g., preferring suboptimal option A or preferring optimal option B), it can be irrational and incoherent to be in both states at once (cf. Goldman 2012: 36, sec 5). For a different reply, see Gert 2007b.

To get appropriate satisficing, merely justifying reasons must interact in a specific way with pro tanto requiring reasons. To get the right kind of interaction, I propose:

Asymmetry: Where O1 is a better option than O2,

- (a) if O2 is not good enough, then the betterness of O1 both pro tanto requires and pro tanto justifies choosing O1 over O2;
- (b) if O2 is good enough, then the betterness of O1 is a merely justifying reason to choose O1 over O2.

For the purposes of the paper, we can say an option is good enough iff, in that option, the agent is well-off, or flourishes. Given this assumption, Asymmetry holds that we are *pro tanto justified* in taking every available improvement to well-being, but we are *pro tanto required* to take only those improvements that we need to be well-off. There is, therefore, an asymmetry in what our well-being pro-tanto justifies and what it pro-tanto requires: it pro-tanto justifies far more choices than it pro tanto requires.¹⁷

If Asymmetry is true, then Necessity is false. Necessity is the claim that when O1 is better than O2, one can appropriately choose O2 only if one has a countervailing consideration. In contrast, Asymmetry claims that when O2 is good enough, one's reason to choose O1 over O2 is merely justifying. Merely justifying reasons can be appropriately ignored. Hence, one can have a merely justifying reason to prefer O1

¹⁷ Although nothing hangs on how generous we are in using the labels “satisficing” or “satisficing theory” (recall the introduction), here is a brief explanation of why Asymmetry should count as satisficing theory and, insofar as a theory allows for appropriate motivated submaximization, it probably should not count as a satisficing theory. It is generally agreed that, to count as “satisficing theory,” a theory must make sense of this idea: choosing a suboptimal option can be appropriate *because the option is good enough* (cf. Henden 2007: 347; van Roojen 2004: 169-70; Weber 2004: 98). Asymmetry makes sense of this idea. The good pro tanto demands that we pursue the good up to a certain point, and this point sets the threshold for the good enough. When I choose an option that meets or surpasses this threshold, the good pro tanto justifies my choice and does not pro tanto require that I choose an alternative option. In the absence of countervailing considerations, my choice is appropriate. Had the option failed to be good enough, the good would have pro tanto required that I choose an alternative. In the absence of countervailing considerations, my choice would have been *inappropriate*. Hence, Asymmetry makes sense of the relevant idea, namely that choosing a suboptimal option can be appropriate because the option is good enough.

van Roojen (2004) and Weber (2004) argue that their defenses of motivated submaximization deserve the label “satisficing theory,” precisely because appropriate motivated submaximization can make sense of the relevant idea. For example, van Roojen (2004: 170) holds that the constraint against torture justifies enduring some loss of the good but not a “calamity”. He infers that his view makes sense of the relevant idea. I’m skeptical of this inference. What he makes sense of is the idea that one can choose a suboptimal option *because the countervailing consideration is strong enough*. Suppose that there are potential countervailing considerations of various strengths. The constraint against lying, say, justifies a loss of up to 10 units whereas the constraint against torture justifies a loss of up to 100 units. One can say, of course, that giving up 50 units of the good for the sake of telling the truth involves choosing an option that is *not* “good enough,” whereas giving up 50 units to avoid torture involves choosing an option that is “good enough”. But here what counts as *being good enough* is determined by the strength of the relevant countervailing consideration: the constraint against torture is strong enough to justify that loss, but the constraint against lying isn’t strong enough. Hence, in appropriate motivated submaximization it’s not clear that *being good enough* plays a genuine, non-redundant explanatory role in making the choice appropriate.

In sum, the contrast is as follows. Insofar as a choice is appropriate satisficing, it is appropriate *because the chosen option is good enough*. Insofar as a choice is appropriate motivated submaximization, it is appropriate *because the relevant countervailing consideration is strong enough*.

over O2 and appropriately choose O2 even in the absence of countervailing considerations.¹⁸

If Asymmetry does capture the central theoretical claim of those who hold that satisficing can be appropriate, then the failure to adequately defend Necessity is even more serious than I first suggested. For then, the objectors' assumption of Necessity just is the assumption that satisficing can't be appropriate. So understood, the Insufficient Reason Objection is arguably a textbook case of begging the question, of assuming that which one aims to prove.

6. Satisficing, Reasons, and Countervailing Considerations

Bradley (2006: 105) will complain that to allow satisficing to be appropriate is to permit "doing a suboptimal act [that] is completely gratuitous". That is, "there is just no reason not to do what is best" (105). It is left implicit that, without such a reason, not choosing the best is inappropriate. Call this the **No Reason Objection**. My response to this objection will be quick; the main work is to show that the resulting account of appropriate satisficing remains distinct from appropriate motivated submaximization.

6.1. Countervailing Considerations

First a quibble: I haven't claimed that an agent can appropriately perform a suboptimal *act* or make a suboptimal *choice*. I claimed that an agent can choose a suboptimal *option*. I think it can be misleading to refer to acts and choices as suboptimal or best (cf. Gert 2000: 237-9, 2007a).

More to the point: on my view, when one satisfices appropriately, there *is* a reason to reject (i.e. not choose) the best. On the conception of reasons I'm working with, R is a reason for choosing option O iff R pro tanto justifies choosing O.¹⁹ The goodness of 9 million units of well-being is a very strong reason—a very strong pro tanto justification—to choose any option that involves that much well-being. And it can't (pro tanto) justify choosing any such option unless it (pro tanto) justifies rejecting all of the other available alternatives. When I can choose between an option with 9 million units and one with 9 million and 1, the goodness of the suboptimal option is a reason that justifies forgoing the best. Appropriate satisficing, then, involves reasons that justify choosing the (mere) good enough *and* justify rejecting the best. Hence, Bradley's objection fails.

At this point, the distinction between appropriate satisficing and appropriate motivated submaximization may seem obscure. Let's clarify it. Recall that satisficing is a type of *unmotivated* submaximization. To say that some submaximization is

¹⁸ Vallentyne (2006: 24) and Russell (2013) endorse something like this response. I think Gert may be committed to something like Asymmetry, but he doesn't explicitly endorse it. Vallentyne cites Kamm (1996, ch 8) as endorsing his version. While Kamm (230-1) mentions Kagan's discussion of something like Asymmetry, she doesn't seem to endorse it. What she endorses is the idea that our autonomy can generate a personal prerogative—a merely justifying reason to prefer one's own projects—that countervails considerations of well-being. As mentioned previously, such a view endorses the existence of merely justifying reasons without thereby endorsing the appropriateness of satisficing.

¹⁹ If you think that a consideration can justify without being a reason, then what really matters is that we have a consideration that justifies not choosing the best. If that consideration fails to count as a reason, then we don't need a reason to reject the best.

unmotivated is to say that it isn't motivated *by countervailing considerations* (sec 3). The key question, then, is this: if appropriate satisficing involves a reason to reject the best, then why doesn't this reason countervail our reason to choose the best? Countervailing considerations operate *independently of* and *against* what they countervail.²⁰ When one satisfices appropriately, one's reason to reject the best is, in the relevant senses, neither independent of nor against one's reason to choose the best.

A reason R1 is *against* choosing option O1 only if, in that situation, R1 does not justify choosing O1. While the acquisition of 9 million units of well-being can justify choosing a suboptimal option (and so justify not choosing the best), it also can justify choosing the best. After all, the best option also exemplifies 9 million units of well-being—and one more unit besides.

There is nothing fishy here. Any theory that allows there to be Buridan's ass cases—cases in which an agent's overall reason justifies choosing either of two options—also allows there to be cases in which the same reason justifies both choosing and rejecting an option. Consider, for example, a crude sort of utilitarianism. On such a view, the appropriateness of a choice is determined entirely by the amount of total well-being contained in one's options and it does not matter how that well-being is distributed across people. Suppose that I have only three options. They are comparable in every respect except as follows. In O1, I benefit no one. In O2, I give Joe 50 units of well-being. In O3, I give Jack 50 units of well-being. Presumably, I can appropriately choose either O2 or O3. Remember that, on this crude utilitarianism, the distinction between persons does not matter. The exact same reason that justifies me in choosing O2—that it leads to an increase of 50 units of total well-being—also justifies me in choosing O3. Since this reason justifies choosing O2, it justifies me in rejecting O3. Nonetheless, since this reason also justifies choosing O3, it also justifies me in rejecting O2. Buridan's ass cases suggest, then, that a reason can justify both choosing and rejecting the very same option (while also justifying both choosing and rejecting an alternative).

If reason R2 countervails reason R1, then there is a minimal sense in which R2 must be independent of R1. Reason R2 is *minimally independent* of reason R1 only if R2 is constituted, at least in part, by something that does not constitute reason R1. Let R2 be the 9 million units that justifies choosing the really great suboptimal option and let R1 be the 9 million and 1 units that justifies choosing the best. Given this very minimal conception of independence, R2 is not independent of R1. Indeed, R1 consists in everything R2 consists in and 1 more unit besides. My reason to choose the really great suboptimal option is not minimally independent of my reason to choose the best.

If we reserve *countervailing considerations* for those reasons that have both againstness and minimal independence, then the term allows us to track the following structural differences. Let "R", "R1", and "R2" represent reasons, while "B" and "E" represent distinct options.

²⁰ I inherit these terms from Slote's (1989, esp ch 2) discussion of the moderate agent, but I develop them in my own way.




Figure 1



Figure 2

Figure 1 depicts the structure involved in appropriate satisficing and Buridan’s ass cases. Let’s focus on the former by letting “B” represent the best option and “E” an excellent suboptimal option. (To represent standard Buridan’s ass cases, both B and E need to be optimal options). The same reason justifies both choosing and rejecting each option, B and E. It’s somewhat odd to think of the reason as countervailing itself. It’s more natural to think of it as a reason that grants you some moral and rational freedom: a single reason makes more than one option choiceworthy.

Figure 2 depicts the structure involved in, say, the sort of Kammian motivated submaximization mentioned previously. One has a reason to choose the best option—it is the best after all. Yet, if choosing the best option requires one to pass up significant opportunities to pursue one’s deeply held projects, then those projects provide a reason to choose a suboptimal option. In this case, the bestness of B is a reason that justifies choosing B (and forgoing E), but this reason does *not* justify choosing E. One’s deeply held projects provide an independent (merely justifying) reason that justifies choosing E (and forgoing B), but it does *not* justify choosing B. These reasons are at odds with one another. They each justify choosing exactly one option, and they justify choosing distinct options. By requiring that countervailing considerations have againstness and minimal independence, we have a natural way of tracking the structural difference between Figure 1 and Figure 2. Figure 1 reasons provide moral and rational freedom, but do not countervail. Figure 2 reasons countervail one another (and, depending on the nature of the reasons, may or may not provide moral and rational freedom).

6.2. The Countervailing Considerations of Motivated Submaximization

We’ve just seen one difference between appropriate motivated submaximization and appropriate satisficing. The former involves a reason that countervails one’s reason to choose the best, i.e., a reason that is against choosing the best and is minimally independent of one’s reason to choose the best. The latter does not involve such reasons. There is a further difference between the two, a difference concerning the kind(s) of countervailing considerations needed for appropriate motivated submaximization.

Suppose that I can bring Albert 50 units of well-being if I choose A, but I can bring Barry 100 units if I choose B. In this case, I presumably have a reason to choose A that is countervailed by my reason to choose B. My reason to choose B is the 100 units it contains. This reason is against choosing A (because A only has 50 units), and it is minimally independent of my reason to choose A (because it consists in 50 more units than my reason to choose A). Nonetheless, my reason to choose A and B are the same kind of reason: they both boil down to the promotion of well-being (or whatever good is at issue). One is stronger than the other, but neither is independent of the promotion of the good.

In appropriate motivated submaximization, however, one’s countervailing consideration for rejecting the best must be independent of the promotion of well-being (or whatever good determines the ranking of one’s options). Suppose that options are ranked solely by the degree of well-being they contain, so that the best option is the one that contains the highest

degree. This degree of well-being provides a reason to choose the best option. Since no smaller degree of well-being is minimally independent of the highest degree, no reason grounded in the promotion of well-being can countervail one's reason to choose the best. Hence, appropriate motivated submaximization involves countervailing considerations that are at least partly independent of the promotion of well-being (or whatever good determines the ranking of one's options).

Let's return to the difference between appropriate motivated submaximization and appropriate satisficing. Both share a common core. They are committed to:

Minimal Teleology: there is some good G, such that the pro tanto appropriateness of choosing an option is determined, at least in part, by the option's degree of G and the degree of G had by the alternatives.

In other words, both appropriate motivated submaximization and appropriate satisficing require that some reasons are grounded in the promotion of G. They nonetheless require two distinct mechanisms for justifying suboptimal choice. Defenders of motivated submaximization reject:

Monogamy: there is some good G, such that the pro tanto appropriateness of choosing an option is determined *solely* by the option's degree of G and the degree of G had by the alternatives.

In the language of reasons, Monogamy claims that there is only one fundamental kind of reason, the kind concerned with the promotion of G. Defenders of motivated submaximization claim that the appropriateness of one's choice partly depends on something independent of G, e.g., the respect of rights. They posit, in other words, another fundamental kind of reason that can compete with or condition or constrain reasons provided by the promotion of G. When we have these competing reasons—these countervailing considerations—they can make it appropriate to choose less than the best with respect to G.

The defender of satisficing need not reject Monogamy. He can say that the appropriateness of a choice depends entirely on the degree of G exemplified by each option. He offers a distinct mechanism, Asymmetry, for making it possible to appropriately choose a suboptimal option. Even if G is the only thing that matters for the appropriateness of a choice, Asymmetry says that not all degrees of G matter in the same way. There are limits on what the good pro tanto demands of us. While it (pro tanto) justifies the best, it (pro tanto) demands the mere good enough.

Appropriate submaximization requires Monogamy to be false but is neutral on Asymmetry. Appropriate satisficing requires Asymmetry to be true but is neutral on Monogamy. These are significant structural differences. And these differences exist even though appropriate satisficing involves a reason to reject the best.

Bradley's No Reason Objection to appropriate satisficing contends that when one satisfices, one has no reason to reject the best. I have argued that this contention is false. The goodness of a good enough suboptimal option provides a reason to choose that suboptimal option. To be a reason that justifies choosing a suboptimal option is to be a reason that justifies rejecting all other alternatives, including the best. Appropriate satisficing, then, involves a reason to reject the best. Appropriate satisficing does not collapse into appropriate motivated submaximization for two reasons. First, the latter involves countervailing considerations and the former does not. When one appropriately satisfices, one's reason to reject the best is neither against choosing the best, nor is minimally independent of one's reason to choose the best. This is no mere verbal gymnastics. By reserving the term

countervailing consideration for reasons that involve againstness and minimal independence, we are able to track interesting structural differences (recall Figure 1 and 2). Second, appropriate motivated submaximization and appropriate satisficing require distinct mechanisms to justify suboptimal choice. Appropriate motivated submaximization requires the existence of reasons that aren't grounded in the promotion of the good. In contrast, appropriate satisficing requires there to be limits on what the good pro tanto demands of us.

7. Satisficing and Means-End Incoherence

The final objection to satisficing is that choosing a suboptimal option in the absence of countervailing considerations—e.g., as a way of optimizing some other, incommensurable good—must involve means-end incoherence (e.g., Byron 1998: 67-9; Dreier 2004: 152; Richardson 2004: 106-9). Recall the case in which I can choose for my life to have 9 million and 1 units of well-being, there are no countervailing considerations, but I choose for it to have 9 million units instead. If my ultimate aim in life is to make my life as well off as I can, my choice frustrates my ultimate aim. This sort of means-end incoherence is widely thought to be irrational. Byron (1998: 69) defines *satisficing* as choosing means to one's end that are merely good enough and not the best.²¹ I agree that, in the absence of countervailing considerations, satisficing *in Byron's sense* necessarily involves means-end incoherence.

Yet satisficing, as I defined it above, does not involve means-end incoherence. If my ultimate aim is to get at least at least 9 million units of well-being for myself, my choice for less than the best (9 million rather than 9 million and 1) coheres with my ultimate aim, it gets me what I really want. Satisficing understood in terms of aims might be necessarily inappropriate for various reasons, but means-end incoherence is not one of them. (To prevent potential confusion: if my ultimate aim is to get at least 9 million units of well-being, there's no incoherence in choosing the optimal option, assumed to be 9 million and 1. Such a choice is a way of getting at least 9 million units and so coheres with my ultimate aim; however, it doesn't count as satisficing because satisficing requires choosing a suboptimal option.)

Indeed, if means-end incoherence is a problem for anything at all, it's a problem for motivated submaximization. In general, it's incoherent to aim at the optimum and then choose the suboptimal, but that's exactly what motivated submaximization involves. So how can motivated submaximization be appropriate? The idea is that the countervailing consideration (e.g., the lack of an optimal option) makes it coherent to choose a suboptimal option even though one aims at the optimum. Or, alternatively, the countervailing consideration prevents the incoherence from counting against the appropriateness of choosing an option. I doubt means-end coherence is genuinely a problem for motivated submaximization. My point is that there is one respect in which satisficing is better off than motivated submaximization: means-end incoherence is a potential problem only for motivated submaximization.

Conclusion

Motivated submaximization and (radical) satisficing are importantly different ways of choosing a suboptimal option.²² To *submaximize with motivation* is to aim at the most good one can get but then choose a suboptimal option because of a countervailing consideration. If such a choice can be appropriate, then Monogamy is false, i.e., it's false that the

²¹ Slote sometimes expresses sympathy with satisficing in Byron's sense (e.g., 1989: 151).

²² In my 2016, I show that this distinction has significant implications for the philosophy of religion.

appropriateness of one's choice is solely a function of how good one's options are. To *satisfice* is to choose a suboptimal option because the option realizes one's aim purely for its own sake at the (mere) good enough. Appropriate satisficing is compatible with Monogamy, but it requires that there are limits on what the good (pro tanto) demands of us. While it (pro tanto) justifies the best, it only (pro tanto) demands the mere good enough.

When satisficing is distinguished from motivated submaximization, it's clear that the existing literature leaves (radical) satisficing in a very bad state. Arguments for the appropriateness of satisficing tend to be better construed as arguments for the appropriateness of motivated submaximization. Or, at best, these arguments show only that either satisficing or motivated submaximization can be appropriate. Furthermore, the existing literature leaves unanswered the three most prominent objections to satisficing.

This paper leaves satisficing in a much better state by answering these objections. The Insufficient Reason Objection fails, because its proponents never bother to defend its key premise, Necessity. The No Reason Objection fails, because it falsely claims that appropriate satisficing does not involve a reason to reject the best. The Incoherence Objection fails, because satisficing needn't involve choosing means that frustrate one's ends. Indeed, if the Incoherence Objection is a problem for anything at all, it is a problem for motivated submaximization.

This paper identified the theoretical commitments of appropriate (radical) satisficing. It showed that these commitments are coherent and worth serious consideration. It did not, however, show that we should endorse those theoretical commitments. That's a task for another day.²³

References

- Bradley, Ben. 2006. "Against Satisficing Consequentialism." *Utilitas* 18: 97-108.
- Byron, Michael. (ed.). 2004a. *Satisficing and Maximizing: Moral Theorists on Practical Reason*. New York: Cambridge University Press.
- _____. 2004b. "Introduction." In Michael Byron 2004a, 1-13.
- _____. 1998. "Satisficing and Optimality." *Ethics* 109 (1): 67-93.
- Cottingham, John. 1986. "Partiality, Favouritism, and Morality." *The Philosophical Quarterly* 36: 357-73.
- Dreier, James. 2004. "Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't." In Michael Byron 2004a, 131-154.
- Gert, Josh. 2016. "The Distinction between Justifying and Requiring: Nothing to Fear." In Errol Lord and Barry McGuire (eds.). *Weighing Reasons* (Oxford: Oxford University Press), 157-172.
- _____. 2007a. "Normative Strength and the Balance of Reasons." *Philosophical Review* 116: 533-62.
- _____. 2007b. "Reply to Tenenbaum." *Canadian Journal of Philosophy* 37: 463-76.

²³ This paper was improved thanks to the feedback of Chris Freiman, Josh Gert, Lanny Goldman, Jonah Goldwater, Aaron Griffith, Matt Haug, Noah Lemos, Tucker McKinney, and the audience at *Fordham University*, not to mention the hard work of anonymous referees. Especial thanks goes to Hamish Russell. Many of the early ideas for this paper were developed while directing his excellent BA Honours dissertation on satisficing. The College of William and Mary also has my gratitude for supporting this paper with a *Faculty Summer Research Grant*. I'm lucky to have an institution and colleagues that are so supportive.

- _____. 2000. "Rationality, Morality, and Purely Justificatory Reasons." *American Philosophical Quarterly* 37: 227-43.
- Goldman, Alan H. 2012. "Response to Gert on Practical Reason." *Journal of Ethics* 16: 35-7.
- Greenspan, Patricia. 2009. "Resting Content: Sensible Satisficing." *American Philosophical Quarterly* 46: 305-317.
- Henden, Edmund. 2007. "Is Genuine Satisficing Rational?" *Ethical Theory and Moral Practice* 10 (4): 339-352.
- Hurka, Thomas. 1990. "Two Kinds of Satisficing." *Philosophical Studies* 59 (1): 107-111.
- Hurka, Thomas and Esther Shubert. 2012. "Permissions To Do Less Than the Best: A Moving Band." *Oxford Studies in Normative Ethics* 2: 1-27.
- Kamm, Frances. 1996. *Morality, Mortality II: Rights, Duties, Status*. New York: Oxford University Press.
- Langtry, Bruce. 2008. *God, the Best, and Evil*. Oxford: Oxford University Press.
- Narveson, Jan. 2004. "Maxificing: Life on a Budget; or, If you would Maximize, Then Satisfice!" In Michael Byron 2004a, 59-70
- Pettit, Philip. 1984. "Satisficing Consequentialism." *Proceedings of the Aristotelian Society* 58: 165-176.
- Pettit, Philip and Geoffrey Brennan. 1986. "Restrictive Consequentialism." *Australasian Journal of Philosophy* 64: 438-455.
- Pollock, John. 1983. "How Do You Maximize Expectation Value?" *Nous* 17: 409-21.
- Richardson, Henry S. 2004. "Satisficing: Not Good enough." In Michael Byron 2004a, 106-30.
- Rogers, Jason. 2010. "In Defense of a Version of Satisficing Consequentialism." *Utilitas* 22: 198-221.
- Russell, Hamish. 2013. "Optimising, Satisficing, and Normative Reasons." B.A. Honours Dissertation. *University of Auckland*.
- Scheffler, Samuel, 2010. "Morality and Reasonable Partiality." In Feltham, Brian, and John Cottingham. *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*, New York: Oxford University Press, 98–130.
- _____. 1982. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. New York: Oxford University Press.
- Schmidtz, David. 1995. *Rational Choice and Moral Agency*. Princeton, NJ: Princeton University Press.
- _____. 2004. "Satisficing as a Humanly Rational Strategy." In Byron 2004a, 30-58.
- Slote, Michael. 1989. *Beyond Optimizing*. Cambridge, MA: Harvard University Press.
- _____. 1984. "Satisficing Consequentialism." *Proceedings of the Aristotelian Society, Supplementary Volume* 58: 139-63.
- Swanton, Christine. 1993. "Satisficing and Virtue." *Journal of Philosophy* 90: 33-48.
- Tenenbaum, Sergio. 2007. "Brute Requirements." *Canadian Journal of Philosophy* 37: 153-71.
- Tucker, Chris. 2016. "Satisficing and Motivated Submaximization (in the Philosophy of Religion)." *Philosophy and Phenomenological Research* 93: 127-43.
- Vallentyne, Peter. 2006. "Against Maximizing Act Consequentialism." In James Dreier (ed.) *Contemporary Debates in Moral Theory*. Malden: Blackwell Publishing, 21-37.

van Roojen, Mark. 2004. "The Plausibility of Satisficing and the Role of Good in Ordinary Thought." In Michael Byron 2004a: 155-175.

Weber, Michael. 2004. "A New Defense of Satisficing." In Byron 2004a, 77-105.