

The Baltic International Yearbook of
Cognition, Logic and Communication

October 2012
pages 1-29

Volume 7: *Morality and the Cognitive Sciences*
DOI: 10.4148/biyclc.v7i0.1773

MICHAEL BISHOP
Florida State University

THE NETWORK THEORY OF WELL-BEING: AN INTRODUCTION

ABSTRACT: In this paper, I propose a novel approach to investigating the nature of well-being and a new theory about well-being. The approach is integrative and naturalistic. It holds that a theory of well-being should account for two different classes of evidence—our commonsense judgments about well-being and the science of well-being (i.e., positive psychology). The network theory holds that a person is in the state of well-being if she instantiates a homeostatically clustered network of feelings, emotions, attitudes, behaviors, traits, and interactions with the world that tends to have a relatively high number of states that feel good, that lead to states that feel good, or that are valued by the agent or her culture.

Change your life and you might make it better or you might make it worse. A change that makes your life go better for you promotes your well-being; a change that makes your life go worse for you diminishes your well-being. This is intuitive enough. But what is well-being? A good way to begin to think about it is with an exercise. How would you explain that a person has a high degree of well-being without actually using the word ‘well-being’ or its synonyms? If you aren’t already corrupted by a philosophical theory, you might offer a thumbnail sketch something like this: “Felicity is in a happy and fulfilling committed relationship, she has close and caring friends, she keeps fit by playing

racquetball, a sport she enjoys, and her professional life is both successful and satisfying.” Most people’s theoretically innocent description of someone with a high degree of well-being will include both “objective” facts about the person (e.g., facts about her relationships, her activities, her professional life) as well as “subjective” facts about the person (e.g., facts about her commitments to and her feelings about kith and kin, facts about her finding certain activities enjoyable and satisfying). The objective and subjective facts we appeal to when explaining Felicity’s—or anyone’s—degree of well-being include:

1. positive feelings, moods, emotions (e.g., joy, contentment),
2. positive attitudes (e.g., optimism, hope, openness to new experiences),
3. positive traits (e.g., friendliness, curiosity, perseverance), and
4. successful interactions with the world (e.g., strong relationships, professional accomplishment, fulfilling hobbies or projects).

So far, so good. But how does this ramshackle set of facts fit into a coherent whole? How are we supposed to unite these various objective and subjective facts into a coherent theory of well-being? The answer I propose is simple: We don’t have to. The world has already joined them together with causal bonds. Think of the above elements of someone’s well-being as nodes in a complex causal network or web. Every node in that network is causally connected to some of the other nodes—it fosters some, and is fostered by others. I call these *positive causal networks*.

The idea behind positive causal networks is that the features of someone’s life we appeal to in explaining her well-being - her committed relationship, friendships, exercise regimens, professional successes, her confidence and sense of mastery, her *joie de vivre*, friendliness, moxie and adventurousness, her curiosity, hope and optimism—are not an accidental conglomeration of happy facts. A person’s well-being has a causal structure—its components are causally bound together. For example, pick out any “happy fact” that is a component of Felicity’s well-being, say, her professional success. This success is fostered by many other factors we appeal to in describing Felicity’s well-being—her curiosity, moxie, optimism, and confidence, her exercise regimen,




Figure 1: Professional Success as a Node in a Positive Causal Network

her social support. But Felicity's professional success is also a cause of some of those "happy facts." Her professional success bolsters her income, her optimism, her confidence and the strength of at least some of her relationships. So Felicity's professional success is a node in a positive causal network. What is true about Felicity's professional success is also true of many of the components of her well-being. Each is a node in a causal network of some combination of positive feelings, positive attitudes, positive traits or successful interactions with the world.

Notice that some states are both cause and effect of Felicity's professional success (her optimism, confidence, and social support). What this suggests is that Felicity's well-being consists of some positive cyclical processes. Felicity's professional success leads her to acquire, maintain or strengthen other positive features of her person; and in turn these positive features help foster her professional success; and so on. These sorts of positive cycles are plausibly associated with many other components of people's well-being. For example, Felicity's optimism helps her overcome challenges and makes her more successful socially and professionally, and having success tends to bolster Felicity's optimism (Seligman 1990). Felicity's friendships and committed relationship provide her with various kinds of material and psychological support, which help to make Felicity more trusting, more extraverted, and more generous, and these traits in turn make Felicity a better friend and partner, which tend to strengthen her friendships and relationship (Fredrickson 2001). Felicity's exercise regimen gives her more strength, energy and positive emotions, which help her, in various ways, to continue her exercise regimen. And so on.

The case for the network theory of well-being (NT) rests on three theses. First, positive causal networks exist. Second, well-being can be understood in terms of positive causal networks. In particular:

(NT) A person's well-being is a function of the strength of her positive causal networks and of her positive causal network (PCN) fragments.

And third, the case for NT is that it organizes, unifies and explains a much wider range of evidence than alternative theories. My goal in this paper is to provide a very rough sketch of a defense of these three theses. A full defense would require a much larger canvass.

1. POSITIVE CAUSAL NETWORKS IN PSYCHOLOGY

The idea that positive causal networks exist is neither radical nor original. Psychologists have identified many instances of positive causal networks, though not under that description, and speculated about their general nature (see, e.g., Lyubomirsky et al. (2005)). The reason this is important, besides giving credit where credit is due, is that most of the job of arguing for the existence of positive causal networks has already been done. My task is to merely clarify the posit and argue that it plays a more significant role in the psychological literature than some might have thought.

Perhaps the most sustained case for the existence and importance of positive causal networks is Barbara Fredrickson's articulation and defense of the Broaden and Build Hypothesis. According to Fredrickson, positive moods and emotions tend to broaden a person's "thought-action repertoire, widening the array of the thoughts and actions that come to mind" (2001, 220). As a result of this broadened thought-action repertoire, the person is more effectively able to build durable physical, social, intellectual and psychological resources "that can be drawn on later in other contexts and in other emotional states" (1998, 307). These resources are durable in the sense that they last much longer than the emotion. Fredrickson's Broaden and Build Hypothesis sets down the following causal schema for a multiply realizable positive causal chain:

Positive affect → Broadened thought-action repertoires → Increased resources.

Add to this the plausible speculation (for which there is quite a bit of empirical support) that having greater social, psychological, material,




Figure 2: Fredrickson's Broaden and Build Cycle

and intellectual resources tends to promote success in ways that foster positive affect. Now what we have is a general schema of an important class of positive causal networks.

A person with a high degree of well-being is in a positive rut or groove—she is enmeshed in a positive causal web involving positive feelings, attitudes, behaviors, traits and successful interactions with the world. This view appeals to commonsense: well-being is being in a “success breeds success” cycle. Of course, it would be just as accurate and incomplete, though considerably less pithy, to say that well-being is being in a “positive feeling breeds positive feeling” cycle, or a “positive attitude breeds positive attitude” cycle, or a “positive trait breeds positive trait” cycle. With this intuitive understanding of positive causal networks in hand, let's focus in more detail on their nature.

2. AN UGLY ACCOUNT OF POSITIVE CAUSAL NETWORKS

Positive causal networks have a commonsense currency. We think that success breeds success, the rich get richer, and good things tend to happen to people who are positive and upbeat. There is a natural temptation for philosophers to try to propose a classical account of positive causal networks: an account framed in terms of singly necessary and jointly sufficient conditions; an account that uses clear and precise

terms; and an account that captures our commonsense, intuitive understanding of the posit. Although I would like to offer a beautiful classical account of positive causal networks, insisting on such an account at this stage of our investigation is a mistake. We are trying to account for a scientific posit, not a commonsense one. We must, of course, *begin* with our commonsense ideas in trying to make sense of the positive causal networks. Where else would we begin? But the ultimate goal is not to capture our commonsense ideas but to improve them. Indeed, relatively young scientific posits are seldom neatly and accurately characterized. Even when science delivers a classical account of a posit like water or lightning, it comes after considerable empirical investigation. Competent investigations into the nature of scientific categories do not typically begin with great conceptual lucidity. It is a mistake to insist on greater conceptual clarity than the subject matter currently affords.

I will propose a modest empirical account of positive causal networks. Rather than explain a posit in terms of singly necessary and jointly sufficient conditions, a modest empirical account identifies characteristics that allow us to reliably identify a posit and distinguish it from other scientific posits. A modest empirical account is explicitly provisional. It can be effectively used to identify the posit. As we learn more about it, we can develop more stylish accounts. Positive causal networks can be identified in terms of three characteristic features.

1. Positive causal networks are made up of an agent's feelings, emotions, attitudes, behaviors and traits, and interactions with the world.
2. Positive causal networks are homeostatic property clusters: A family of properties that tend to co-occur because “[e]ither the presence of some of the properties... tends to favor the presence of the others, or there are underlying mechanisms or processes which tend to maintain the presence of the” property cluster (Boyd 1989, p. 16). As with many homeostatic property clusters, positive causal networks will have “borderline” cases in which there is no fact of the matter about whether something is or is not a positive causal network, and they might be without an essence—there is no property or mechanism that occurs in all instances of the cluster (Boyd 1989, p. 16–17).

Now let's consider the question of what makes a causal network—a

homeostatically clustered set of feelings, emotions, attitudes, behaviors, traits, and interactions—a *positive* causal network.

3. A homeostatically clustered network of feelings, emotions, attitudes, behaviors, traits, and interactions with the world is *positive* if it consists of relatively more of the following sorts of states:
 - a. psychological states that feel good—that have a positive hedonic tone;
 - b. states (psychological or not) that when present in this network tend to bring about psychological states that have a positive hedonic tone;
 - c. states that the agent values;
 - d. states that the agent's culture values.

Causal networks that are positive tend to have relatively more states that feel good, that lead to states that feel good, or that are valued by the agent or her culture. Given our current state of knowledge, I take it that we are not in a position to offer any very informative account what it is for a psychological state to have a “positive hedonic tone.” The best we can do is repeat the wisdom of Louis Armstrong, who, when pressed to define jazz, is reputed to have said, “Man, if you gotta ask, you’ll never know.” We can, of course, point to stereotypical examples of experiences with positive hedonic tone, such as the physical pleasures (e.g., sexual, gastronomic), the feelings involved with positive social interactions (love, close friendship) and aesthetic experiences (e.g., listening to a symphony, looking at great art). But I will proceed on the assumption that positive hedonic tone is a basic, unanalyzable notion. A call to explain it is a cry for help, not a request for serious philosophical inquiry.

Positive causal networks can have varying degrees of strength. Some change, C, to a person's positive causal networks can strengthen or weaken those networks. C strengthens a person's positive causal network if C makes it more robust—it makes the network more able to persist for a longer time and in a wider range of plausible environments. C weakens a person's positive causal network if C makes it less robust—it makes the network less able to persist for a longer time and in a wider range of plausible environments. The intuitive idea, here, is that positive causal networks are homeostatic systems, like living organisms or running engines. Changes strengthen the system when

they make the system tougher, sturdier, more durable, harder to extinguish. Changes weaken the system when they make it more delicate, less durable, easier to extinguish. The factors that alter the strength (or robustness) of a positive causal network are discoverable only empirically. There is a lot to say on this matter. But for our purposes, I will leave it at this: In general, there are two ways to strengthen or weaken a person's positive causal networks: (i) by changing the intensity of the states that compose the networks or (ii) by changing the size of those networks (i.e., by increasing or decreasing the number of states that make them up). Usually, but not always, an increase in intensity (e.g., more happiness, better relationships, more success at work) strengthens a positive causal network. And usually, but not always, an increase in the size of a network (e.g., new friends, newly achieved success at work) strengthens a positive causal network.

The modest account of positive causal networks I have offered here will win no beauty contests. Good thing, too. The account merely aims to mark off one sort of thing from other sorts of thing. My fondest wish is that it should someday soon be replaced by something far superior.

3. POSITIVE CAUSAL NETWORK FRAGMENTS

Joy's life is going exceptionally well. She instantiates a very robust set of positive causal networks. Now consider a series of Joys (Joy', Joy'', etc.) each with a missing link: each subsequent Joy is missing one more link from Joy's original positive causal networks. Joy' lacks one link, Joy'' lacks that link and one more, and so on. The lives of any two contiguous Joys would be going (almost) equally well. Eventually, there will be intermediate cases such that there is no fact of the matter about whether these Joys instantiate a positive causal network. Such borderline cases are to be expected. Then at some further point down the line, Joy (or rather, Joy''') would clearly not instantiate a positive causal network. But she would still possess fragments of a positive causal network—states that *could* be part of a positive causal network. And intuitively, the presence of these positive causal network *fragments* would make Joy's life better than their absence. To properly account for well-being, we need to account for these positive causal network fragments (or PCN fragments).

The need for PCN fragments arises from the fact that it is possible for a person to be better or worse off even if he does not instantiate any positive causal networks. Even if George fails to instantiate any positive causal networks, his life might go better being entertained by Harold Lloyd's *Safety Last* than ruminating over his latest misery. George enjoying this great silent movie would be a fragment of a positive causal network. A state (or set of states) is a fragment of a positive causal network just in case it is a type of state that *could* be a significant link in a positive causal network for that person, keeping relatively constant the sort of person he is (i.e., his personality, his goals and his general dispositions). Take all the plausible positive causal networks George *might* be in. Given his temperament and abilities, for example, he might instantiate various networks involving his profession, his social relationships and his hobbies. The types of states causally implicated in these networks are all PCN fragments for George. The states that make up these fragments—the pleasant experiences or moods, the successes, the exercises of talents or skills—they all make life better for George.

A set of states might be a PCN fragment for one person but not another. For example, suppose Arthur is a misanthrope, unable to instantiate any positive causal networks involving friendship. It's not that he *can* thrive in close friendships but prefers not to; and it's not that he is friendless against his wishes. Rather, Arthur is constitutionally incapable of thriving in relationships with other people. In that case, his having a friend, by itself, would not be a PCN fragment. Keep in mind that given Arthur's misanthropy, his having a friend would not engender in him the typical feelings of camaraderie, solidarity and support it would in the rest of us. Arthur's friend might, of course, help him to act or feel or be in ways that *are* PCN fragments for Arthur. And so his friend might make Arthur's life better indirectly. But the mere fact that Arthur has a friend does not make his life better. Another less dramatic example of this sort of phenomenon is that successfully engaging in some act of mild daredeviltry might be a PCN fragment for Daring Dan but not for Cautious Charlie.

Fragments of positive causal networks can have varying degrees of strength. Some change, C, to a person's PCN fragment can strengthen or weaken that fragment. What makes for degrees of strength for PCN

fragments? This is a tricky matter that deserves more space than I have here. But in general, a PCN fragment is stronger insofar as it consists of a greater number of states or those states are more intense (e.g., more positive affect, more success, etc.). There is one significant qualification: A PCN fragment is weakened by any change that makes it more difficult for the person to instantiate a positive causal network. This exception is necessary because too much of an otherwise good thing can sometimes be bad for you. For example, in environments where one might face genuine risks and dangers, greater degrees of positive affect, hope and optimism can undermine well-being by engendering a lack of caution (e.g., engaging in risky behavior, not taking preventative measures). Optimism among older folks "predicts depression in the wake of stressful events. Perhaps extreme optimism among the elderly is unrealistic, and the occurrence of something terrible can devastate the optimistic older individual..." (Peterson & Seligman 2004, p. 577).

So the network theory of well-being holds that a person's well-being is a function of the strength of her positive causal networks and of her PCN fragments. Other things being equal, a person's well-being is promoted with any increase to the strength of her positive causal networks or fragments thereof; and a person's well-being is diminished with any decrease to the strength of her positive causal networks or fragments thereof. Now let's turn to the case for NT.

4. WHAT DO WE WANT FROM A THEORY OF WELL-BEING?

What we want from a theory of well-being is straightforward enough: We want a theory that accurately and perspicuously describes the nature of well-being. But how are we to achieve this goal? What evidence should we consider? And how should we proceed from that evidence? The traditional approach to the study of well-being embraces the Descriptive Adequacy condition: a successful theory must capture our commonsense judgments about well-being. James Griffin states that "the notion we are after is the ordinary notion of 'well-being'" (1986, p. 10). In L.W. Sumner's discussion of what we want from a theory of welfare, which he takes to be "more or less the same" as well-being (1996, p. 1), he concisely articulates the first part of the traditional

approach:

[T]he best theory about the nature of [well-being] is the one which is most faithful to our ordinary concept and our ordinary experience. That experience is given by what we think or feel or know about well-being, both our own and that of others. The data which a candidate theory must fit, therefore, consist of the prodigious variety of our preanalytic convictions (1996, p. 10–11).

The best theory of well-being is the one that “makes the best sense of [our preanalytic] convictions” (1996, p. 11). The degree of fit between a theory of well-being and our pretheoretic convictions is “a function of the extent to which the truth conditions [the theory] offers can support and systematize our intuitive assessments.” While Sumner does not argue that “descriptive adequacy” is the sole requirement that the correct theory of well-being must satisfy, it is clearly the most important—it is “the basic test” (1996, p. 10). Valerie Tiberius also embraces the Descriptive Adequacy condition:

Formal analyses [which provide an account of the nature of well-being] are to be evaluated on the basis of how well they accommodate our uses of the concept in question and how well they fit with our ordinary experience. In other words, formal accounts of well-being are evaluated primarily in terms of their descriptive adequacy. The most descriptively adequate account of well-being is the one that is most faithful to our pre-philosophical convictions about well-being (2005, p. 299).

On the traditional approach, the way to evaluate a theory is primarily in terms of whether it is “faithful” to our commonsense judgments about well-being.

The fundamental problem with the traditional approach is that it is epistemically over-optimistic. It assumes that our commonsense intuitions about the nature of well-being are so epistemically special that it is reasonable that they should serve as the primary base of evidence for a theory of well-being. The best way to appreciate the epistemological immodesty of the traditional approach is in terms of the diversity challenge: Different people have different commonsense judgments about

well-being. It is not clear how the traditional approach is to choose among them. To make the case for diversity in people’s commonsense judgments, we do not need to conduct experiments or travel to exotic locales. There is a robust diversity in commonsense judgments about well-being among Western philosophers who are experts on the subject.

The Experience Machine. Consider two people who have exactly the same experiences, but one is genuinely engaged with the world and the other is prone in a laboratory with a machine feeding electrical impulses into her brain (Nozick 1974). Do the two people with exactly the same experiences have the same degree of well-being? Some philosophers think they don’t (Nozick 1974) while others, including many hedonists, think they do (Crisp 2006).

Remote Desires. We have desires that extend in time and space far beyond our ken. Examples of remote desires include the desire for a posthumous good reputation, the desire for a stranger to flourish, the desire for some distant future scenario (e.g., functional jet packs by the 24th century), or some quirky desire whose satisfaction is epistemically inaccessible (e.g., a prime number of atoms in the universe) (Parfit 1984; Griffin 1986; Kagan 1998). Does satisfaction of these remote desires promote our well-being? Insofar as these remote desires do not impinge upon our experience, classical hedonists think their satisfaction cannot affect well-being. Among desire theorists, there is a range of opinions. Mark Lukas argues that satisfaction of every actual desire, including remote desires, promotes well-being, although he seems to readily admit that this requires that one “embrace the absurdity and simply deny the intuition that some desires are irrelevant to well-being” (Lukas 2010, p. 21). Mark Overvold argues that the only desires whose satisfaction promote a person’s well-being are those whose satisfaction logically require her existence; and so the satisfaction of remote desires does not promote well-being (1982). Other philosophers, however, argue that satisfaction or frustration of posthumous desires can affect a

person's well-being (Brandt 1979; Kavka 1986; Portmore 2007). James Griffin distinguishes between informed satisfied desires that can and cannot count towards a person's well-being as follows: "What counts for me, therefore, is what enters my life with no doing from me, what I bring into my life, and what I do with my life" (1986, p. 22). I interpret this to mean that as long as a remote informed desire is properly connected to one's life plan, its satisfaction promotes the person's well-being. Griffin's restriction rules out some remote desires (e.g., the jet pack or prime number of atoms desires) but not all of them (e.g., the desire for a posthumous good reputation).

For the traditional approach, these cases are meant to elicit a commonsense judgment that is supposed to form part of the evidential base for a theory of well-being. But philosophers's commonsense judgments about these cases are incompatible. How are such disagreements to be resolved? This is a tricky predicament because philosophers are adept at fitting their commonsense judgments into coherent theories of well-being. Each philosopher begins with her own idiosyncratic evidential foundations. And they proceed to build an assortment of clever, interesting and sometimes beautiful theories. But without some principled way to decide which evidential foundations are the right ones, whose commonsense judgments are correct, the traditional approach runs the risk of congealing into a sterile stalemate. Given the current state of the debate, one can be forgiven for thinking this is something more than a risk.

I propose an approach to the study of the nature of well-being that begins with the assumption that well-being is a real state or condition. It is a condition that scientists can learn about and that laypeople talk about, perhaps imperfectly, when we talk about well-being. The assumption that well-being is a real state or condition demands a certain epistemic modesty about our commonsense judgments about well-being. On the one hand, our commonsense judgments cannot be completely off-base. They must track well-being well enough that we can successfully talk about it. On the other hand, there is no reason to begin our investigation by supposing that our reflective, pretheoretical judgments about well-being are completely coherent and accu-

rate. As our investigations proceed, of course, we might have reason to revise our epistemic modesty. Perhaps we will discover that our commonsense judgments are extremely accurate descriptions of well-being. Perhaps we will discover that most, but not all, of our commonsense judgments are reasonably accurate descriptions of well-being. A rather pessimistic possibility is that no condition in the world answers very well to our commonsense judgments. In that case, perhaps we conclude that well-being is a failed posit, like caloric or aether, or perhaps we conclude that it is a real condition but it is not the sort of condition that can be empirically studied. A strength of approaching the study of well-being by assuming that it is a real condition is that we need not prejudge which of these outcomes is true. Perhaps our commonsense judgments will turn out to be highly accurate. But perhaps not. Deep confidence or deep pessimism about our commonsense well-being judgments should be the *result* of our investigation into the nature of well-being, not the *starting point* of that investigation.

The base of evidence a theory of well-being must account for extends far beyond the confines of commonsense. A good theory of well-being will capture at least reasonably well our commonsense judgments about well-being. But because our commonsense judgments might not be perfectly accurate, this cannot be the only goal of theorizing. A theory of well-being should also be able to unify, organize, and make sense of the science of well-being. While this approach is unconventional, it does not repudiate the traditional approach. It does not deny that a theory of well-being should account for our commonsense judgments. But it asks for significantly more. A theory of well-being should organize, unify and make sense of positive psychology. A theory of well-being deserves our allegiance insofar as it does the best job of accounting for the *entirety* of the evidence. This philosophical approach to the study of scientific posits is common in philosophy of science, particularly philosophy of psychology and mind. For example, longstanding debates about the existence and nature of belief, emotions and consciousness take as evidence both commonsense lore and scientific findings (e.g., Stich (1983); Griffiths (1998); Block (2007)). In *Knowledge and its Place in Nature* (2002), Hilary Kornblith uses essentially this approach to articulate and defend a theory of knowledge. Kornblith's work is an important precedent since he applied this ap-

proach to knowledge, which, like well-being, is taken to be a normative category.

5. HOW NT MAKES SENSE OF POSITIVE PSYCHOLOGY

Most scientific disciplines or subdisciplines can be characterized in ways that are pithy and specific.

- Cytology is the study of the structure, composition and function of cells and their parts.
- Kinematics is a branch of mechanics that studies motion (the continuous change of position).
- Biochemistry is the study of the chemical substances and processes that occur in living organisms.
- Cognitive psychology is the study of mental processes such as perception, memory and reasoning.

These characterizations provide a clear sense of what these disciplines are about by identifying (putatively) real categories in nature that are their object of study—cells, motion, chemical substances, living organisms, perception, reasoning. Positive psychology is different. Some characterizations of positive psychology are pithy and vague, while others are specific and prolix. But none are like the descriptions of cytology, kinematics, biochemistry or cognitive psychology—pithy and specific. Here are some pithy but vague characterizations of positive psychology.

Positive psychology is the scientific study of what goes right in life, from birth to death and at all stops in between (Petersen 2006, p. 4).

Positive psychology aims to help people live and flourish rather than merely to exist (Keyes & Haidt 2003, p. 3).

The label of *positive psychology* represents those efforts of professionals to help people optimize human functioning by acknowledging strengths as well as deficiencies, and environmental resources in addition to stressors (Wright & Lopez 2005, p. 42).

Compared to the characterizations of other scientific disciplines, these are desperately vague and unclear. Other characterizations of positive psychology have the opposite problem. They are full of specifics but are so prolix that they seem to include everything but the kitchen sink.

The field of positive psychology at the subjective level is about valued subjective experiences: well-being, contentment, and satisfaction (in the past); hope and optimism (for the future); and flow and happiness (in the present). At the individual level, it is about positive individual traits: the capacity for love and vocation, courage, interpersonal skill, aesthetic sensibility, perseverance, forgiveness, originality, future mindedness, spirituality, high talent, and wisdom. At the group level, it is about the civic virtues and the institutions that move individuals toward better citizenship: responsibility, nurturance, altruism, civility, moderation, tolerance, and work ethic (Seligman & Csikszentmihalyi 2000, p. 5).

This list, as long as it is, is not an exhaustive itemization of what positive psychology investigates. It is a laundry list of “for examples”, not a coherent description of a scientific discipline. All this suggests a lack of clarity concerning what positive psychology is really about.

It is noteworthy that all these different characterizations of positive psychology seem true. This fact can be explained by the hypothesis that positive psychology is the study of positive causal networks. If positive psychology is the study of positive causal networks, then it is also the study of “what goes right in life”, and it *does* have the capacity to “help people... flourish” and “optimize human functioning...” And the prolix characterization is accurate insofar as it touches on the three subjective elements of positive causal networks: positive feelings and emotions (contentment, satisfaction, happiness), positive attitudes (hope, optimism), and positive traits (courage, perseverance, originality, altruism, tolerance, civility). It omits the fourth element of positive causal networks, successful interactions with the world. But some above average success in the world typically accompanies any long stretch of the subjective items cited.

The hypothesis that positive psychology is the study of positive causal networks makes sense of how practitioners of positive psychol-




Figure 3: A Positive Causal Network

ogy characterize their field of expertise. It also helps to organize the large body of research that flies under the banner of positive psychology. Consider what positive psychology says about friendship. It is most natural to represent this visually rather than with the written word. Most of the items that appear in figure 3 are familiar to commonsense, except perhaps for Positive Affect. I take this link to represent both a relatively stable disposition to have positive experiences as well as the occurrence of transient positive experiences. While this amalgam is conceptually unlovely, it helps keep the chart simple. This chart represents an idealized, incomplete empirical hypothesis. Take any proposed causal connection posited in the chart. There is at least some evidence for that connection, but there are reasonable complaints one might raise - one might object to how the states were measured, to the experimental arrangements, to whether there is good evidence for causal connections among the states rather than mere correlations, etc. But my contention is that there is enough evidence to suppose that something like this causal network exists even if some of its pieces should not survive further investigation. The chart is also rad-

ically incomplete. The arrows typically represent indirect causal connections that can be mediated in many different ways. For example, the “Success with Projects → Positive Affect” chain is well-established. But one would be right to point out that it is typically *perceived* success rather than *actual* success that is causally relevant to positive affect. Success that is not perceived as success won’t lead to positive affect; and lack of success that is perceived as success can lead to positive affect. So the chart should perhaps include the intervening link.

Success → Perceived Success → Positive Affect

But neither of those chains is direct or unmediated. The “Perceived Success → Positive Affect” chain does not always obtain - depressive episodes often survive perceived successes, after all. Besides including incomplete oversimplifications, this chart does not aim to represent any real positive causal networks, which are messy and complicated. It is best understood as an idealized and incomplete instance of a positive causal network.

The visual representation of this positive causal network makes clear that it is, in fact, a *network*, involving links bound together with many causal connections. The network consists of many positive cycles—connections that loop back onto the same types of states. Begin at any node and a sequence of causal connections will take you to any other node. As a result—and this can be lost in any linear, written description—there is no compulsory starting point. There is no state we *must* privilege as the most important in the network. Mind you, there might *be* some states that are of particular importance to this network. But these questions go to the heart of positive psychology as the study of the *dynamics* of positive causal networks rather than their *structure*. Dynamical questions concern what factors scuttle, inhibit, maintain, promote or establish positive causal networks. For example, there may be some link that is typically or practically necessary for the above friendship network to remain in operation. If that link is weakened or removed, the entire network is crippled or destroyed. Just by looking at the chart, it seems as though *Positive Affect* is a good candidate for such a state. But this node covers a wide range of states, some might be important, others not. Furthermore, the chart is incomplete. For all we know, there might be a set of nodes not included in this chart that would keep a (somewhat modified) version of this network going even

if positive affect were largely absent. There may also be some link that is typically or practically sufficient to establish or maintain this network in operation. In reasonably favorable circumstances, it might be that the presence of some link will kick start the friendship network. Commonsense suggests that in reasonably propitious environments, *Others Judging One More Positively* together with *Extraversion* might be sufficient to start this process, whereas *Positive Affect* by itself isn't. And it is natural to wonder whether some link in the friendship network is more important than others in one of these ways: Is some link usually or practically necessary for the operation of the entire network? Does some link or set of links, in friendly environments, typically establish such a network? These are good questions, questions naturally raised by thinking about positive psychology as the study of the dynamics of positive causal networks. But they are empirical questions to be settled by competent empirical research.

Let's focus on some of the parts of the friendship network. When positive affect is induced in the laboratory, studies suggest that it will tend to make you more sociable and friendly (3 → 1). For example, you are more likely to start a conversation with a stranger (Isen 1970). And in that conversation, you are more likely to offer intimate self-disclosures (Cunningham 1988). What's more, you will be more generous and positive in your judgments and interpretations of other people. For example, after conducting a simulated job interview, positive affect subjects rate interviewees more highly and are more likely to "hire" them (Baron 1987). So people who are in a good mood and feeling happy will tend to be friendlier, more open and more generous toward other people. But this is in the laboratory. What about in the real world? People high in positive affect tend to judge their interactions with others to be more pleasant and enjoyable than people low in positive affect. For example, happier people are more likely to express a desire to be friends with or work on a project with a new acquaintance, and they are more likely to judge the person to be "kind, self-assured, open, tolerant, warm" (Lyubomirsky & Tucker 1998, p. 179).

The good vibes happy people send out to others are reciprocated in spades (3 → 6). People high in self-reported positive affect are more favorably judged by the people they interact with as well as by third parties who view videotapes of their interactions (Berry & Hansen

1996, p. 800). A review of the literature reports that happy people are judged to be better looking, more competent and intelligent, friendlier and more assertive, more moral "and even more likely to go to heaven." The friends and family of happier people judge them to be more "socially skilled (e.g., more articulate and well mannered), better public speakers, self-confident, and assertive, and as having more close friends, a strong romantic relationship, and more family support" (Lyubomirsky, King & Diener 2005, p.827), see text for citations).

Our discussion of friendship has focused considerably on positive affect. But to emphasize the point made above that positive causal networks have no a priori compulsory starting points, we could have started our discussion with the personality trait *extraversion*. In a longitudinal study, Costa and McCrae found that extraversion (e.g., sociability, vigor, social involvement) predicts positive affect and life satisfaction 10 years later (1980, p. 675) (1b → 3). In a meta-analysis, DeNeve & Cooper (1998) argue that the extraversion-happiness correlation is quite strong; in fact, they argue for the existence of a positive cycle involving positive affect, friendly personality traits, and successful relationships.

Positive affect is not tied solely to Extraversion. Rather, positive affect stems primarily from our connections with others, both in terms of the quantity of relationships (Extraversion) as well as the quality of relationships (Agreeableness)... [R]elationship type personality traits foster better relationships. However, they appear to provide another bonus to the holder; they also facilitate the experience of positive affect... (220-221).

At the risk of beating a long dead horse, I am not imposing positive causal networks on the psychology from the armchair. It's right there. Costa and McCrae argue that positive affect brings about extraversion, which brings about more and stronger relationships, which brings about positive affect... And so on.

6. NT AND COMMONSENSE

The case for NT relies on arguing that NT does about as well as competitors at capturing our commonsense well-being judgments and it is

far superior at making sense of positive psychology. To show that NT captures commonsense reasonably well, my plan is to argue that NT can explain both what is intuitively right and intuitively wrong with a pair of well-known theories of well-being, hedonism and informed desire theory (IDT). It is important to keep in mind the logic of the case for NT. The objections to hedonism and IDT I will consider will be very familiar to proponents of those theories. And defenders of these theories will have smart things to say about these objections. But this is irrelevant. The case for NT is *not* that it captures the intuitions of the proponent of IDT better than IDT. *It doesn't*. NT does not capture the hedonist's intuitions better than hedonism, and it does not capture the Aristotelian's intuitions better than Aristotle's theory. In fact, NT doesn't even capture my own commonsense judgments perfectly. The right conclusion to draw from all this is that arguing for a view of well-being on the grounds of commonsense is a recipe for dissensus and deadlock. If I get to judge your theory in terms of whether it captures *my* intuitions, then in the face of my settled insistence that my theory does a better job than your theory at capturing *my* intuitions, you don't have much more to say. We are at an impasse. But if the main goal of your theory is to capture a wide range of evidence of which my commonsense judgments are but a small part, then in the face of my firm insistence that your theory doesn't capture my intuitions, you have a straightforward and compelling reply. You can rationally defeat my resistance by pointing to the fallibility of commonsense judgment and to the far greater explanatory power of your theory. Of course, I might insist upon the epistemic sublimity of my commonsense judgments and refuse to budge. But that would be my problem, not yours.

6.1. Hedonism

Hedonism is the thesis that well-being is a function of the balance of positively valenced experience (e.g., enjoyment, pleasure, happiness) over negatively valenced experience (e.g., suffering, pain). It is the thesis that "what is good for any individual is the enjoyable experience in her life, what is bad is the suffering in that life, and the life best for an individual is that with the greatest balance of enjoyment over suffering" (Crisp 2006, p. 622). Hedonism is an explanatory theory. It explains why close relationships or professional success contribute

to a person's well-being. What makes those things good for a person, the reason they contribute to her well-being, is that they bring about a robustly favorable hedonic balance of enjoyment over suffering.

NT accounts for the intuitive power of hedonism. If NT is true, hedonism is approximately true. Positive causal networks consist of states with a robustly favorable hedonic balance of enjoyment over suffering. And so according to NT, people with more well-being tend to have more (net) positive experiences. Hedonism and NT will yield a considerable amount of overlap in their particular judgments about people's well-being. They diverge in their explanatory ambitions. Suppose Susan begins some sort of treatment regimen that bolsters her well-being. The hedonist explains this in terms of the overall improvement in the hedonic tone of Susan's subjective experiences. She now has more net pleasure than she did before. But for NT, the treatment regime brings about a significant increase in Susan's well-being because it makes more robust her positive causal networks. It might change the intensity of certain parts of her positive causal networks. For example, she is a bit less industrious, which permits her to appreciate other successful aspects of her life, or she is a bit more patient with family and friends, which strengthens her close relationships. Or perhaps the treatment regimen adds new elements to her positive causal networks. For example, she is now more at peace with her place in life and more optimistic about the future which leads her to form new friendships and rekindle old ones. These changes make Susan's positive causal networks more robust—they are more resilient to life's occasional knocks. For NT, the reason Susan's treatment regime increases her well-being is that it brings about stable changes to her life. It bolsters the robustness of the positive causal networks that make up her well-being. The hedonic zing is part of the story, and a crucial part of most stories, but it is only a part. NT unearths and makes explicit the causal structure and the dynamics of well-being, and it explains well-being in terms of factors that are causally implicated in its perpetuation. From the perspective of NT, the explanations offered by the hedonist for why some aspect of Susan's life contributes to her well-being are usually only partially true—on the right track but incomplete.

Some will object that I have unfairly saddled hedonism with a short-sightedness that is contrary to the view. Suppose the short-term

hedonic zing of Susan's treatment regime brings about dispositions of behavior or of mind that lead to further pleasant experiences. The hedonist would rightly insist on counting these longer-term hedonic consequences of the treatment regime as part of the explanation for its effects on Susan's well-being. This is a fair point. But from the perspective of NT, the problem with the hedonist's explanation is not that it is short-sighted. It is that the explanation, no matter how far sighted, is inevitably partial. *It ignores the causal structure, stability and dynamics of well-being because it focuses exclusively on one part of that structure.* The hedonic tone of a person's life is typically a good indicator of the strength of the positive networks that make up her well-being. In fact, it might be the best single indicator there is. And so hedonism is a reasonable approximation of the truth about the nature of well-being. But it is not the whole truth.

To appreciate how well NT conforms to our everyday intuitions about well-being, let's consider a standard objection many philosophers take to be a serious problem for hedonism. This objection targets any mental state view of well-being, any view that takes a person's well-being to be entirely a function of her mental states. It relies on Nozick's experience machine thought experiment.

Suppose there were an experience machine that would give you any experience you desired. Superduper neuropsychologists could stimulate your brain so that you would think and feel you were writing a great novel, or making a friend, or reading an interesting book. All the time you would be floating in a tank, with electrodes attached to your brain... Would you plug in? (Nozick 1974, p. 42–43).

We can sharpen the example by supposing that Richard and Anthony have exactly the same positive experiences, mostly positive, except that Richard is hooked up to the experience machine while Anthony is genuinely engaged with the world. Hedonism, like any mental state theory of well-being, implies that the experience machine doesn't matter to a person's level of well-being. As long as they're having exactly the same experiences and exactly the same mental states, hedonism yields the result that Richard and Anthony have exactly the same levels of well-being. And the conventional wisdom among philosophers is that this is wrong. Anthony has a higher level of well-being than Richard. My

argument here is not that hedonism is false because it fails to accord with many philosophers's intuitions. Instead, the experience machine objection is relevant because what philosophers have said and thought about well-being is part of the evidence a theory of well-being should explain. And in his defense of hedonism, Roger Crisp describes the current conventional philosophical wisdom:

Hedonism has a distinguished philosophical history... In the twentieth century, however, hedonism became significantly less popular... [W]hile hedonism was down, Robert Nozick dealt it a near-fatal blow with his famous example of the experience machine. The result has been that these days hedonism receives little philosophical attention, and students are warned off it early on in their studies, often with a reference to Nozick (2006, p. 619–20).

This is roughly what we would expect if NT were true and philosophers over time have been getting closer to the truth about the nature of well-being. While there is much more to say on this topic, my provisional conclusion is that compared to hedonism, NT does reasonably well at capturing people's commonsense well-being judgments.

6.2. Informed Desire Theory

The basic idea behind desire theories is that a person is better off insofar as she gets what she wants. More carefully, IDT holds that a person's well-being is a function of the satisfaction of her informed desires. A desire is satisfied (or fulfilled) when the content of the desire comes about. So a person's well-being is a function of whether certain states in the world obtain, and in many cases those states are not subjective mental states. For example, when I desire that my child learn to swim, what satisfies that desire is not that I believe that he has learned to swim or that he or I enjoy his learning to swim. What satisfies my desire is that my child actually learns to swim. If this is one of my informed desires and it is satisfied, then according to a standard version of IDT this increases my well-being. But not all desires are created equal. My satisfied desire that my children be healthy increases my well-being significantly more than my satisfied desire that my favorite sports team wins this weekend. James Griffin argues that the strength

of a desire is not a matter of intensity or motivational force. At some time, the desire one has with the strongest motivational force might be to smoke or overeat, even though satisfaction of those desires might not increase one's well-being. "If strength [of desire] were interpreted as motivational force, then 'utility' would lose its links with well-being... So to retain the links with well-being, the relevant sense of 'strength' has to be, not motivational force, but rank in a cool preference ordering, an ordering that reflects appreciation of the nature of the objects of desire" (Griffin 1986, p. 15).

A significant challenge for any IDT is to explain what it is for a desire to be informed. The motivation for this caveat is clear: Some of our actual desires can be based on ignorance or misinformation. For example, the satisfaction of my desire to eat the delicious looking banana split or to take the euphoria-inducing drug might not increase my well-being at all, and in fact might undermine it, because the banana split is tainted and because the drug leads to addiction and ruin. Had I been appropriately informed about these matters, and had I used that information in a rational way, that desire would not have survived. But what is it for a desire to be appropriately informed? I think the best answer is given by Peter Railton. He accounts for an individual's "objectified subjective interest" as follows.

Give to an actual individual A unqualified cognitive and imaginative powers, and full factual and nomological information about his physical and psychological constitution, capacities, circumstances, history, and so on. A will have become A+, who has complete and vivid knowledge of himself and his environment, and whose instrumental rationality is in no way defective. We now ask A+ to tell us not what he currently wants, but what he would want his nonidealized self A to want—or, more generally, to seek—were he to find himself in the actual condition and circumstances of A... [W]e may assume there to be a reduction basis for his objectified subjective interests, namely, those facts about A and his circumstances that A+ would combine with his general knowledge in arriving at his views about what he would want to want were he to step into A's shoes (2003, p. 11).

So my informed desires are those desires my idealized self would want me (the non-idealized me) to want were he to be in my shoes.

NT can capture what is intuitively plausible about IDT. If NT is true, then the IDT is approximately true, although strictly false. Satisfying a person's informed desires will typically help to promote the robustness of her positive networks. This isn't an accident. If we focus on people's actual desires, healthy people with a modicum of insight about themselves and how the world works tend to have a host of desires that, if satisfied, would promote the robustness of their positive causal networks. In general, most of us want our relationships to be strong, we want our families to prosper, we want to be good to our friends and we want our friends to be good to us, we want to be happy, healthy, safe and productive. These very general desires will engender particular desires that will vary widely given different people's situations, e.g., you want to land that job, and I want to retire early. The overlap between NT and desire theory is even greater when we restrict ourselves to informed desires—i.e., the desires your idealized self would want for your real self if your idealized self were in your shoes. When a person's particular desires are well-informed—when, for example, retiring early really will have the positive consequences on my life that I think it will—the satisfaction of these desires will tend to strengthen the positive networks that make up a person's well-being. So NT can explain what's intuitively right about IDT.

The most basic objection to IDT is that the connection between informed desires and well-being is a contingent one. Some satisfied desires do nothing for well-being, while others can positively undermine well-being. Remote desires (desires whose satisfaction makes no difference to our experience) are one class of desires that *prima facie* do nothing for a person's well-being. Eric's desire for posthumous recognition (e.g., fame or something more modest, like a fancy carving on his tombstone) might survive full information. But would the satisfaction or non-satisfaction of such posthumous desires affect his well-being? Desires whose satisfaction are epistemically inaccessible and that do not touch on Eric's life (e.g., he desires that there be a prime number of atoms in the universe) seem to be irrelevant to well-being (Parfit 1984; Griffin 1986; Kagan 1998). Rawls gives an example of a person who wants to count the blades of grass on the college green (1971,

p. 432). It seems possible for the satisfaction of this desire to actually undermine well-being. Most philosophers take at least some of these implications to be contrary to commonsense. As a result, proponents of IDT often adopt various embellishments to avoid at least some of these implications.

NT takes these objections to IDT to be largely correct. Take any desire Eric might have—no matter how strong or well-informed that desire might be. It is always an open empirical question how the satisfaction of that desire will affect the positive causal processes that constitute Eric's well-being. The satisfaction of posthumous desires or epistemically inaccessible desires will do nothing for the robustness of his positive causal networks. And so the objection that satisfaction of certain remote desires is irrelevant to a person's well-being is true. What's more, it is possible that the satisfaction of an informed desire might undermine the robustness of a person's positive causal networks. The grass counter might spend his time on a project that brings nothing but physical discomfort, grief and derision. If the satisfaction of that desire undermines the smooth operation of the positive causal networks that make up Eric's well-being, then according to NT it would undermine his well-being.

There is much more to say about NT and commonsense. But this is enough to at least make it plausible to suppose that NT does an adequate job capturing our commonsense judgments about well-being. For now, that is enough.

7. CONCLUSION

This paper offers an approach to the philosophical investigation into well-being that begins with the assumption that well-being is a condition that can be empirically studied. The approach is novel insofar as it has never (as far as I am aware) been applied to the study of well-being. But it is by no means novel to philosophers. It rests on very standard and familiar philosophical views about kinds and how we learn about them, and it has been used by many philosophers to investigate philosophically interesting subjects.

On the assumption that well-being is a real condition, no theory about its nature can win the day simply by providing a coherent ac-

count of someone's commonsense judgments about well-being. Given the diversity of such judgments, commonsense *by itself* is simply not epistemically sturdy enough to adjudicate between competing theories. And so the case for NT is built on a simple strategy: Battle to a draw on the commonsense evidence and win on the scientific evidence. NT is not so contrary to commonsense that it deserves to be forsaken. And it does exceptionally well at unifying and making sense of the scientific evidence. In order for another theory to defeat NT, it is not enough to show that it captures someone's commonsense judgments better than NT. One must show that it captures the *entirety* of the evidence better than NT.¹

Notes

¹For an excellent discussion of the topics raised in this paper, I would like to thank the audience at the 7th International Symposium of Cognition, Logic and Communication: Morality and the Cognitive Sciences in Riga, Latvia. Special thanks to Francesco Orsi for pressing me to appreciate the importance of positive causal network fragments. I would also like to thank audiences at Union College, at the Universidad de Los Andes in Bogota, at the Symposium on Naturalism in Science at Kansas State University, at the conference on Ethical and Social Scientific Perspectives on Well-Being at California State University, Long Beach, and at the Moral Philosophy Research Group Workshop at Washington University, St. Louis

References

- Baron, R. A. 1987. 'Mood interviewer and the evaluation of job candidates'. *Journal of Applied Social Psychology* 17: 911–926.
- Berry, D. S. & Hansen, J. S. 1996. 'Positive affect, negative affect, and social interaction'. *Journal of Personality and Social Psychology* 71, no. 4: 796.
- Block, N. 2007. 'Consciousness, Accessibility and the Mesh between Psychology and Neuroscience'. *Behavioral & Brain Sciences* 30: 481–548.
- Boyd, R. 1989. 'What Realism Implies and What It Does Not'. *Dialectica* 43: 5–29.
- Brandt, R. B. 1979. *A Theory of the Good and the Right*. Clarendon Press: Oxford.
- Costa, P. & McCrae, R. 1980. 'Influence of Extraversion and Neuroticism on Subjective Well-Being: Happy and Unhappy People'. *Journal of Personality and Social Psychology* 38, no. 4: 668–678.
- Crisp, R. 2006. 'Hedonism Reconsidered'. *Philosophy and Phenomenological Research* 73, no. 3: 619–645.
- Cunningham, M. 1988. 'Does Happiness Mean Friendliness? Induced Mood and Heterosexual Self-Disclosure'. *Personality and Social Psychology Bulletin* 14, no. 2: 283–297.
- DeNeve, K. & Cooper, H. 1998. 'The happy personality: A meta-analysis of 137 personality traits and subjective well-being'. *Psychological Bulletin* 124, no. 2: 197–229.

- Fredrickson, B. 1998. 'What Good are Positive Emotions?' *Review of General Psychology* 2 3: 300–319.
- . 2001. 'The Role of Positive Emotions in Positive Psychology: The Broaden-and-Build Theory of Positive Emotions'. *American Psychologist* 56, no. 3: 218–226.
- Griffin, J. 1986. *-Being: Its Meaning, Measurement, and Moral Importance*. Clarendon: Oxford.
- Griffiths, P. 1998. *What Emotions Really Are*. University of Chicago Press.
- Isen, A. 1970. 'Success, Failure, Attention, and Reaction to Others: The Warm Glow of Success'. *Journal of Personality and Social Psychology* 15, no. 4: 294–301.
- Kagan, S. 1998. *Normative Ethics*. Westview Press: Boulder, CO.
- Kavka, G. 1986. *Hobbesean Moral and Political Theory*. Princeton University Press.
- Keyes, C. & Haidt, J. 2003. *Flourishing: Positive Psychology and the Life Well-Lived*. American Psychological Association.
- Kornblith, H. 2002. *Knowledge and Its Place in Nature*. Oxford University Press: New York.
- Lukas, M. 2010. 'Desire Satisfactionism and the Problem of Irrelevant Desires'. *Journal of Ethics and Social Philosophy* 4, no. 2: i.
- Lyubomirsky, S. & Tucker, K. 1998. 'Implications of Individual Differences in Subjective Happiness for Perceiving, Interpreting, and Thinking About Life Events'. *Motivation and Emotion* 22, no. 2: 155–186.
- Lyubomirsky, S., King, L. & Diener, E. 2005. 'The Benefits of Frequent Positive Affect: Does Happiness Lead to Success'. *Psychological Bulletin* 131, no. 6: 803–855.
- Nozick, R. 1974. *Anarchy, State, and Utopia*. Basic Books.
- Overvold, M. 1982. 'Self-Interest and Getting What You Want'. In H. Miller & W. Williams (eds.) 'The Limits of Utilitarianism', 186. University of Minnesota Press.
- Parfit, D. 1984. *Reasons and Persons*. Clarendon: Oxford.
- Peterson, C. 2006. *A Primer in Positive Psychology*. Oxford University Press: New York.
- Peterson, C. & Seligman, M. 2004. *Character strengths and virtues: A handbook and classification*. Washington, DC: American Psychological Association.
- Portmore, D. 2007. 'Welfare and Posthumous Harm'. *American Philosophical Quarterly* 44, no. 1: 27–38.
- Railton, P. 2003. 'Facts and Values'. In *Facts, Values, Norms: Essays Towards a Morality of Consequence*. Cambridge University Press.
- Rawls, J. 1971. *A Theory of Justice*. Harvard University Press.
- Seligman, M. 1990. *Learned Optimism*. Knopf: New York.
- Seligman, M. & Csikszentmihalyi. 2000. 'Positive Psychology: An Introduction'. *American Psychologist* 55, no. 1: 5–14.
- Stich, S. 1983. *From Folk Psychology to Cognitive Science*. MIT Press.
- Sumner, L. W. 1996. *Welfare, Happiness, and Ethics*. Oxford University Press: New York.
- Tiberius, V. 2005. 'Cultural Differences and Philosophical Accounts of Well-Being'. *Journal of Happiness Studies* 5: 293–314.
- Wright, B. & Lopez, S. 2005. 'Widening the Diagnostic Focus: A Case for Including Human Strengths and Environmental Resources'. In C.R. Snyder & Shane J. Lopez (eds.) 'Handbook of Positive Psychology', 26–44. Oxford University Press: New York.