Pluralism and the Levels of Selection in Cancer

by

Joseph Wu

Department of Philosophy
Duke University

Honors thesis submitted in fulfillment of the
Graduation with Distinction for the Philosophy Major

2015

Abstract

The recognition that cancer is an evolutionary process has yielded many insights for understanding the dynamic nature of the disease. But an evolutionary perspective bequeaths a philosophical problem that has largely been ignored by cancer researchers, namely, the levels of selection. In this paper I explain how cancer is an evolutionary process and why the levels of selection issue is central to understanding the origin and existence of cancer. I then clarify the difference between pluralist and monist interpretations of the levels of selection disputes, a philosophical distinction that many scientists are unaware of. Applying James Woodward's theory of causation, I argue for pluralism and show how a genomic and multicellular model of cancer parses the causes differently to provide an equally accurate representation of carcinogenesis. I conclude by suggesting pragmatic advantages of genomic, cellular, and multicellular models of cancer in biomedical research. My analysis illustrates the value of adopting the philosophical thesis of pluralism in the practice of science.

Pluralism and the Levels of Selection in Cancer

1. Introduction

Cancer has been described as an evolutionary process since the seminal work of Cairns and Nowell in the 1970s (Cairns, 1975; Nowell, 1976). Under this view, a neoplasm can be viewed as a genetically heterogeneous population of individual cells that resulted from and are under the influence of natural selection and genetic drift (Figure 1). Genetic abnormalities with different, heritable effects on the fitness of cells lead to selection for 'renegade' cells, which successively acquire the functional adaptations characteristic of cancer cells (Greaves and Maley, 2012). As the value of evolutionary biology in cancer research is gaining traction, the naïve notion that cancer is a static entity with a single cause is being replaced by the recognition of its complex and dynamic nature (Gatenby, 2009; Greaves and Maley, 2012).

But an evolutionary perspective of cancer bequeaths a philosophical problem. Natural selection operates on many levels in biology. The orthodox story of cancer is told at the cellular level. Research is directed at understanding the role of the tumor microenvironment, the effects of neighboring cells, resource availability, and predation by immune cells. But there is an alternative selective story. The story of selection can also be told at the level of the genome, and as we shall see, sometimes at the level of groups of cells. The latter stories have been largely absent from the scientific literature. Why is this so? Can a genomic or multicellular perspective complement or, possibly, replace the cellular perspective? What is the relationship between the levels of selection and cancer?

This paper argues that the evolution of cancer is best captured by a pluralist approach to the levels of selection (Waters, 1991; Waters, 2005). Pluralists contend that in some cases, multiple alternative models can each truthfully represent an evolutionary process. It is the view

that a true model "is not necessarily *uniquely* true," and that in these instances no model is the single correct way to represent the causal story involved (Waters, 2005, pg. 312). There are multiple equally adequate accounts of evolutionary events. I think pluralism is right. I also think that cancer is an evolutionary process through which a pluralistic approach would be best illustrated and vindicated, at least in the sense that pluralists recommend "practicing biologists take advantage of the full range of strategies for representing the workings of selection" (Sterelny and Kitcher, 1988). In a field aimed at therapeutic improvement, a pluralistic approach to studying the evolution of cancer is not just epistemologically and ontologically significant. It can also have empirical and clinical implications.

Pluralism about the levels of selection also raises a second question, as Waters (2005) notes: "If some situations can be individually accounted for by a plurality of true models, as pluralists contend, could scientists nevertheless have rational grounds for choosing some true models over others?" Waters claims that there are pragmatic reasons for choosing between different models. Yet he does not provide any specific examples of how this might be carried out—that is, how alternative models of an evolutionary process can yield different practical pay offs depending on which model is chosen. In this paper I aim to address this deficiency by extending Waters' argument and detailing the benefits of genomic, cellular, and multicellular models of cancer evolution. I suggest that each of these true models of the evolutionary process highlight important aspects of cancer evolution that the other models overlook or cannot account for.

The discussion unfolds as follows: in Section 2 I provide a brief overview of why cancer is an evolutionary process and how cancer is currently modeled at the cellular level. Following that, in Section 3, I explain how the levels of selection issue is central to understanding the origin

and existence of cancer. In Section 4, I clarify and defend the version of pluralism with respect to the levels of selection advanced in this paper. In Section 5, I apply James Woodward's theory of causation to show how a genomic model can accurately represent cancer as an evolutionary process. I also discuss recent studies in the cancer literature that suggest the appropriateness of multicellular models. Finally, in Section 6, I outline pragmatic advantages of genomic, cellular, and multicellular models in studying the evolution of cancer.

Glossary

Neoplasm	Abnormal mass of tissue resulting when cells divide more than they should or do not die when they should. Also called a tumor.
Missense mutation	Single nucleotide change that results in a codon for a different amino acid
Nonsense mutation	Single nucleotide change that results in a premature stop codon
Frameshift mutation	Insertions or deletions of nucleotides that alter the reading frame of the codons and result in translation of a different protein
Deletion mutation	Genetic aberrations in which part of a DNA sequence is missing
Subclones	Descendant of a mutant arising in the previous generation
Angiogenesis	Formation of new blood vessels
Metastasis	Spread of cancer from where it arose to distant locations in the body
Somatic	Cells of the body that are not the germ-line cells which give rise to ovum or sperm
Lymphocytes	White blood cell that determines the specificity of the immune response to foreign substances
Visceral Organs	Internal organs of the body, in particular those within the chest or abdomen

Figure 1. Glossary of cancer terminology

2. The Evolutionary Nature of Cancer

Evolution by natural selection occurs when there is phenotypic variation, heritability, and differential fitness causally related to differential adaptedness. The process of carcinogenesis satisfies these prerequisites. Cancer arises from random DNA alterations, including missense, nonsense, frameshift, and deletion mutations (Bignell *et al.*, 2010; Youn and Simon, 2011; Greenman *et al.*, 2006). This random and purposeless production of genetic diversity leads to an assortment of different cellular phenotypes, providing the heritable variation on which natural

selection can operate. Some mutations are advantageous to the cell while others are neutral or harmful. Beneficial mutations increase the fitness of 'renegade' cells, enhancing their propensity for survival and reproduction (Greaves and Maley, 2012; Brandon, 1978). As more genetic alterations accumulate, the mutant cells and/or their variant subclones successively acquire the functional capabilities that characterize cancer cells, each of which confers a selective advantage over normal cells. Hanahan and Weinberg (2011) organized these phenotypes into the 'Hallmarks of Cancer,' which include sustaining proliferative signaling, evading growth suppressors, resisting cell death, enabling replicative immortality, inducing angiogenesis, activating invasion and metastasis, reprogramming of energy metabolism, and evading immune destruction. The multistep process of carcinogenesis can be rationalized as the acquisition of functional traits that enable incipient cancer cells to achieve replicative success and, eventually, immortality in a tumor microenvironment.[1]

Consider how one trait, the ability to evade immune destruction, increases the expected fitness of the cancer cell. Recent studies have demonstrated that immune surveillance plays a key role in tumor eradication. When mice were genetically engineered to be deficient in various aspects of the immune system, it was observed that tumors occurred more frequently and/or proliferated more rapidly in the immunodeficient mice relative to the immunocompetent controls. Specifically, deficiencies in $CD8^+$ cytotoxic T lymphocytes, $CD4^+$ T_h1 helper T cells, or natural killer cells each resulted in observable increases in tumor incidence. These results suggest that the immune system contributes to the eradication of cancer cells (Teng *et al.*, 2008; Kim *et al.*, 2007). What this means for the cancer cell, however, is that the ability to avoid immune

[1] Immortality is defined here as biological or cellular immortality. In biology, cells are immortal if they have an infinite growth capacity—that is, if these cells can be cultured indefinitely in a lab. However, these 'immortal' cells can still be killed, for instance with chemical agents. Thus, they are not immortal in a colloquial sense—they do not, strictly speaking, live forever.

destruction is a selective advantage over both other cancer cells and other normal cells. The immune system is part of an environment selecting against cancer cells, and only those cells that possess the ability to circumvent external cellular predation can succeed in it. Hence, evading immune destruction is an adaptation that causally leads to increased fitness of the cancer cell.

So cancer development is an evolutionary process in which mutant cells or their descendent subclones gain fitness advantages over other cells in a particular microenvironment from the heritable variation that arose via random mutations. This useful framework for understanding the nature of cancer progression, however, has continued to be overlooked in the scientific literature, with only about 1% of cancer papers on therapeutic resistance or relapse using evolutionary terms in the abstracts since the 1980s (Aktipis *et al.*, 2011). This is problematic for at least two reasons. First, an evolutionary framework of cancer biology is explanatorily superior to a framework that focuses solely on the molecular mechanisms mediating carcinogenesis. This is because an evolutionary framework can provide both proximate *and* ultimate explanations for how and why cancers arise. For instance, when asked why a normal epithelial cell transformed into a malignant one, a molecular biologist might reference the activation of particular oncogenes that deregulate the checkpoints in the cell cycle (a proximate explanation), while an evolutionary cancer biologist would also cite this cell cycle deregulation but consider it an adaptation that increases the chances of reproductive success for the cancer cell (thus providing both a proximate and ultimate explanation). Second, an evolutionary framework in cancer biology is empirically relevant in the clinic. Tumors are heterogeneous populations of cells. Cancer therapies apply strong selective pressures on these cells, and in cases where not all of the malignant cells are eradicated, these therapies select for the resistant populations that can survive, continue to proliferate, and later lead to the recurrence

of a tumor that is composed of a new and diverse population of cells resistant to the therapy applied earlier (Figure 2). Thus, adopting an evolutionary framework in cancer treatment is necessary in order to increase predictive ability in the clinic and further improve the methods of therapeutic intervention.

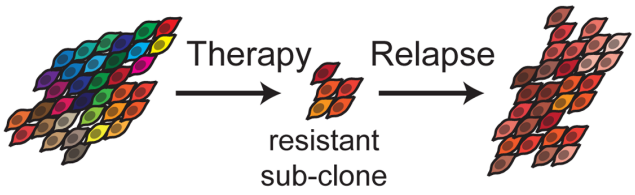

Figure 2. Cancer therapy selects for resistant cells in a heterogeneous population. Relapse occurs when a resistant subclone survives and continues to proliferate to form a new and diverse population of resistant cells. [Reprinted from Aktipis *et al.*, 2011.]

To sum, in this section I have explained how cancer is an evolutionary process and why an evolutionary framework in cancer research and treatment is crucial. Note that the account thus far frames cancer as a cellular selection process. This is the orthodox story told in the evolutionary cancer literature. But is there more to the story?

3. Cancer and the Levels of Selection

The prerequisites for natural selection are heritable phenotypic variation that results in differential reproduction. These requirements are abstract. In principle, they can apply to *any* entities that satisfy them. Further, biologists have long recognized that the biosphere is hierarchically organized, with more recent work aiming to characterize the nature of these levels of nestedness (McShea 1996, 1998, 2001). In the biological hierarchy, organisms constitute one level, but there are levels above and below that of the organism, which can also satisfy the prerequisites for natural selection. As explained in the previous section, for instance, there can be

natural selection for cells, a level below the organismic level. Given the abstract character of natural selection and the hierarchical nature of the biosphere, natural selection can occur at many levels of biological organization. Indeed, the formation of the biological hierarchy is itself a product of selection at different levels during the major transitions of evolution (Maynard Smith and Szathmáry, 1995). But what does it mean for selection to occur at a given level? And in a particular evolutionary event, to what extent does selection act at each level? In this section, I show why these levels of selection questions are central to understanding the origin and existence of cancer.[2] For our purposes here, it is worth noting that natural selection occurs at a given level "if entities at that level vary with respect to a character which causally influences fitness" (Okasha, 2006 pg. 27).

Evolutionary cancer researchers have gestured toward the levels of selection in their work. For example, Aktipis and Nesse (2012) note the tension between cell-level selection and organism-level selection as a critical factor in the most important 'major transition' in the history of life—the origin of multicellular organisms (Maynard Smith and Szathmáry, 1995). In particular, as organisms lived longer and increased the number of cells in their bodies, the ability to reduce mutations and suppress cancer became evermore crucial for the survival of the organism (Caulin and Maley, 2011). The tension between organism-level and cell-level selection required the organism to develop powerful mechanisms to suppress evolution at the somatic level in order to be favored by natural selection at the organismic level. Thus, natural selection explains both why there is cancer and why it is not more common amongst multicellular organisms.

[2] There is a separate levels of selection question that concerns identifying the evolutionary processes that led to the formation of the biological hierarchy itself (Griesemer, 2000; Okasha, 2006). As Michod (1999) emphasizes, multicelled organisms did not simply arise from nowhere. While this aspect of the levels of selection question is undoubtedly important, it falls beyond the scope of this paper.

A particularly interesting example that exemplifies how cancer relates to the levels of selection is the Tasmanian devil facial tumor disease (DFTD). DFTD is one of two known cancers that are transmissible between hosts.[3] DFTD is transmitted through biting, and it frequently metastasizes to lymph nodes and visceral organs (Loh *et al.*, 2006). As a result, most Tasmanian Devils that are infected survive no more than 6 months (Hawkins *et al.*, 2006; Lachish *et al.*, 2007). In fact, this cancer is so virulent that if the current rate of disease spread continues, DFTD poses a risk of extinction for the Tasmanian Devils in the next 25 – 35 years (McCallum *et al.*, 2007). Consequently, selection for the infectious cancer is so strong that it may result in the demise of the organismic host necessary for its maintenance and transmission. Despite this, transmission frequency remains high, which suggests that there is little pressure for DFTD to reduce virulence (Lachish *et al.*, 2007; McCallum, 2008). It would be interesting to explore whether DFTD can manipulate their hosts to increase transmission, perhaps by promoting Tasmanian Devil aggression, and whether DFTD clones that have decreased virulence are selected for in the long run by allowing Tasmanian Devils to survive through the mating season (Murchison, 2009).

In DFTD, we have a clear case of natural selection operating in opposing directions at two different levels—an increased fitness of the cancer is correlated with a decreased fitness of the Tasmanian Devil host. But this raises two questions. The first relates to what it means for a cancer cell to have a higher fitness in this scenario. What makes DFTD so intriguing is the fact that it is not bound to the host like most cancers. It is transmissible, so the death of the host does not necessitate the ending of the DFTD cell lineage, for the descendants may have already been passed onto another Tasmanian devil through a bite. Hence, whereas the fitness of most cancers

[3] The other is canine transmissible venereal tumor, which is an infectious genital tumor in dogs that is spread during coitus (Cohen, 1985).

is measured by cellular replication rate during the lifespan of the organism, the fitness of DFTD can also be measured by its ability to leave descendants across organismic generations. A DFTD lineage may be less fit in the first regard while on the face of a Tasmanian devil compared to normal cells because it divides slowly, but nevertheless be more fit in the second regard because it is less virulent than a rapidly dividing lineage and therefore more likely to be transmitted to another host to survive across organismic generations. Any discussion of DFTD fitness, then, will be contingent upon which environmental timeframe is considered, as well as how the environment is conceptualized—does the environment for DFTD include only the face of one Tasmanian devil, or does it include the faces of all Tasmanian devils within a population? The understanding of fitness as a notion specific to a relative timeframe and environment will be essential in later discussion.

The second question raised concerns what, exactly, makes DFTD a levels of selection issue. Recall that selection at a given level occurs when entities possess heritable variation that causally leads to differential reproduction. Intuitively, both DFTD and Tasmanian devils satisfy these prerequisites. There is selection at both the cellular and organismic level, at the same time, and it happens that what is advantageous at one level is disadvantageous at another. The direction of selection is in opposite directions at different levels. But these are empirical matters of fact. If that is all there is to the levels of selection issues, then there does not seem to be much dispute at all. So what is all the fuss about? Indeed, up until now the discussion has focused on largely uncontroversial claims that most, if not all, biologists and philosophers of biology would endorse. For the remainder of the paper, I turn my attention to some of the philosophical issues involved in the levels of selection disputes regarding pluralism, causation, and how our understanding of them influences what is at stake in the scientific literature. Having laid the

groundwork for why the levels of selection are relevant to cancer evolution, my aim from here on out will be to consider how the extensive philosophical literature on the levels of selection can complement and further develop scientific research in evolutionary cancer biology.

4. Pluralism and the Levels of Selection

The philosophical literature on the levels of selection has focused on establishing criteria for identifying which entity or entities are selected for in a given scenario, and why it is that these entities satisfy the prerequisites for evolution by natural selection. In this section, I argue for a certain type of pluralism with respect to the levels of selection. Before I continue, however, a conceptual distinction must be made between 'replicators' and 'interactors,' which are generalizations of the genotype-phenotype distinction. David Hull introduced this helpful terminology, refining Dawkins' distinction between 'replicators' and 'vehicles' (Hull, 1980; Dawkins, 1978). In "Individuality and Selection," Hull notes that Dawkins' notion of a "replicator" interacts with the environment in two ways: to produce copies of itself, and to influence its own survival as well as the survival of its descendants. Hull introduces the term 'interactor' to refer to the latter process. According to Hull, a replicator is "an entity that passes on its structure directly in replication," and an interactor is "an entity that directly interacts as a cohesive whole with its environment in such a way that replication is differential" (Hull, 1980, pg. 318). Hence, selection is "a process in which the differential extinction and proliferation of interactors cause the differential perpetuation of the replicators that produced them" (ibid.).

Traditionally understood, the interactor denotes what entities are actively being selected in the process of natural selection.[4] Note that in the previous section, the story of cancer evolution involved two distinct interactors. On the one hand, there were the cancer cells that had acquired advantageous mutations allowing them to leave more descendants of themselves in a particular microenvironment. On the other hand, there were the Tasmanian Devils, which possessed organismic traits that affected their expected reproductive success in an ecological environment. We have already noted that these were empirical matters of fact. There are two separate selective processes occurring with two interactors at different levels of the biological hierarchy.

But let us put aside the issue of organismic selection for now and focus solely on the evolution of cancer. From section 1 we saw that the current model of cancer evolution in the scientific literature is in cellular terms. It is indisputable that this model is beneficial for understanding and explaining the process of carcinogenesis. Yet is it the only model available? Many scientists and philosophers assume that for a given selection process, there must be a *true* or *real* level of selection, and that this is a fact of matter about the particular evolutionary event. However, there is no *a priori* reason to assume that this is always the case. While some models will misrepresent evolutionary processes, these can be rejected on empirical grounds. In particular cases, alternative models that represent the same evolutionary situation might not disagree about important matters of fact. If this is true, then are there other interactors available besides the cell that can truthfully model cancer evolution?

Pluralism about the levels of selection is the view that in some cases there are multiple equally adequate accounts of evolutionary events that can be modeled at one or more levels of

[4] Robert Brandon (1982) draws a distinction between the *units* of selection question and the *levels* of selection question, where the former refers to the level of replicators and the latter to the level of interactors. I will focus on the level of interactors in this paper; that is, on the levels of selection question and not the units of selection issue.

the biological hierarchy. For example, a situation that is modeled as selection for individuals might also be accurately modeled as selection for alleles. Contrast this with monism, the thesis that for a given situation, there is a uniquely correct model that captures the *true* level(s) of selection. According to monists, in one situation, the *real* level of selection might be the organismic level, while in another situation the *real* level of selection might be the cellular level. In a third scenario there might be genuine selection at both the organismic and cellular levels. Waters (2005) succinctly expounds the distinction as follows: "*Pluralism about levels of selection* is the view that monisms about the targets of selection are false. It says that in *some* situations, an evolutionary process can be modeled correctly as selection for units at one level, and alternatively modeled correctly as selection for units at another level (or multiple levels)."

The version of pluralism I am defending warrants more discussion and clarification given the many ways the term has been used in the literature. One very weak sense of pluralism is the view that selection *can* occur at different levels. This is what Sober and Wilson (2002) refer to as "pluralism concerning what happens in nature" (pg. 529), and it is uncontroversial. The evolution of cancer cells on the faces of Tasmanian Devils and the formation of the biological hierarchy vindicate this pluralism. A more philosophically interesting version of pluralism claims that there is no fact of the matter about the true level(s) of selection. This thesis can be local or global in scope. While a local view would contend that in particular cases there is no objective fact about the true level(s) of selection, a global view would contend that there is never such a fact of matter. Sterelny and Kitcher (1988) and Kitcher (2004) defend global pluralism. On Sterelny and Kitcher's (1988) view, it is always a metaphysical error to assert there is a 'true' level of selection, for there is something to be said about the proposed equivalence between being able to tell the selection story in different ways, with the genic point of view "always available"

(Sterelny and Kitcher, 1988, pg. 359).[5] Later, Kitcher (2004) advances a slightly different view by going so far as to say that, "the unit of selection controversy…has become a philosophical exercise in irrelevant metaphysics." Alternatively, Dugatkin and Reeve (1994) argue for a local pluralism. They believe that pluralism is accurate with respect to a particular class of evolutionary processes, and for reasons specific to that class. The pluralism I am endorsing here is local in scope. It depends only on the idea that *some* evolutionary situations can be modeled accurately by alternative theories, and I will show that this is the case in cancer in the following section. It is worth emphasizing that the pluralism I am advancing has two aspects to it, one concerning causality and the other concerning the possibility of multiple representations. I explore these two aspects in turn to further elaborate the version of pluralism I am defending here.

4.1 *Pluralism about Causality*

While a pluralist might adopt a non-realist view of causation, this is not the pluralism I am arguing for here. On a non-realist account, there is no objective matter whether a trait *causally* influences fitness. For example, Kitcher (2004) says that it does not make sense to inquire about "the *real* locus of causation in selection processes," as it involves a mistaken reification of causal relations (emphasis in original). He writes:

[5] Some philosophers such as Lloyd (2005) have suggested that the arguments for genic pluralism by Sterelny and Kitcher "entail genic reductionism", and that this pluralism implies that the genic view is superior to other views because it is always available. I agree that this charge is problematic for the global pluralism of Sterelny and Kitcher, for it does seem odd to privilege one particular level on a pluralist account. This is one reason why I think their version is incorrect. The local pluralism I advocate does not categorically privilege any particular level, although it does claim that in the evolutionary situations that can be accurately modeled by alternative theories, a particular level can be privileged on pragmatic grounds.

> One can tell all the facts about how genotype and phenotype frequencies change across the generation – including the causal explanation of the changes – without any commitment to a definite level at which selection acts. For example, it makes no difference whether one thinks that a particular allele's production of a specific protein initiates a causal chain that makes it likely that that allele will have an enhanced chance of finding its way into the gene pool or whether one believes that some phenotypic trait raises the probability that an organism will have increased reproductive success (Kitcher, 2004).

This argument seems to me to be wrong. There *is* a difference between the different loci of causation in selection processes, and it *is* certainly reasonable to take a realist stance on statements such as 'X causally affects Y.' As Okasha (2006) explains, "To ask about the 'real locus of causation' is to ask about the hierarchical level or levels at which character differences cause fitness differences…If the explanation for why a given particle character has changed in frequency is that it causally affects the fitness of particles, for example, then selection has acted at the particle level; if not, then not." (pg. 128). So given that we accept a realist view of causation, which is more consistent with biological research than a non-realist stance on causation, Kitcher's claim seems incorrect. It is clearly not possible to provide a causal explanation for why a trait increased fitness 'without any commitment to a definite level at which selection acts.'

But Kitcher's view alludes to another version of pluralism that takes a realist stance on causality, and similarly asserts that the causes involved in evolutionary events are complicated enough that they can be cut up in different ways. This pluralism that I am advancing holds that the causal story at one level is not the complete causal story, and as a result, the causes of a particular evolutionary event can sometimes be described truthfully at different levels. There are multiple causal stories. On this view, Kitcher is right to claim that there are many causal chains that can explain 'how genotype and phenotype frequencies change', but wrong to assert that this

does not commit us to any particular level of selection, for the level at which the causal story is told is the level at which selection acts. As we will see later, Kitcher is also wrong to think 'it makes no difference' which causal story is chosen, as I suggest that this decision can be made on grounds that have practical implications. The basic thesis here is that "the causes in evolutionary situations are sufficiently complicated that they can be described truthfully in multiple ways" (Waters, 2005).[6]

Two points are worth noting here. First, the multiplicity of causal stories leads to a 'causal selection' problem—how do we choose between different causal chains when explaining how genotype and phenotype frequencies change? Biologists typically do not explain an evolutionary event by citing all of the conditions *necessary* for its occurrence, which is what John Stuart Mill called the 'total cause'. Rather, they privilege some causal chains and relegate the other factors to the status of 'background' or 'enabling' conditions. Many philosophers of biology have aimed to provide principled solutions to the causal selection problem (see Franklin-Hall, forthcoming; Waters, 2007; and Strevens, 2008); however, discussion of these attempts falls beyond the scope of this paper. For our purposes here, I wish to only remark that the selection of causes can be considered a pragmatic matter, loosely echoing Mill's argument that causal selection is 'capricious'—"Nothing can better show the absence of any scientific ground for the distinction between the cause of a phenomena and its conditions, than the capricious manner in which we select from among the conditions that which we choose to denominate the cause" (Mill, 1846, pg. 198). Mill's view, I think, is consistent with the version of pluralism

[6] This is an epistemological thesis that "calls for a less strident form of scientific realism (at least with respect to evolutionary biology" (Waters, 2005). One might wonder whether this claim depends on the current state of scientific knowledge, or whether evolutionary phenomenon are so complicated that a single and complete causal story can never be told due to the limitations of the human mind. I am inclined to say that it is the former. In principle at least, I think a being with higher intellect or better tools could identify all of the causal chains and present one comprehensive causal story. But given the current state of science, this complete causal story is currently (albeit not forever) inaccessible to humans.

advanced in this paper, for if the selection of causes is in fact capricious, then it allows the pluralist to model the causes of an evolutionary event in different ways without worrying about whether the model isolates the single, 'real' causal chain.

Second, the pluralism we are interested in will strictly concern causality, not what might be called 'bookkeeping.' These two aspects must be kept distinct. One might, for instance, be a pluralist in the sense that different models can successfully trace and predict changing frequencies (bookkeeping), while be a monist with respect to identifying the underlying causes that change those frequencies. The levels of selection debate turns on whether, at a particular level, there is a *causal* relationship between a trait and an increase in fitness relative to a specific environment. Hence, while success in bookkeeping is undoubtedly important, the issue at hand concerns a deeper kind of success involving the identification of the causes that bring about those changes in frequencies. For the remainder of this paper, I use the terms monism and pluralism in the causal sense only.

4.2 *Pluralism and Alternative Models*

The second aspect of the pluralism I am arguing for here concerns the possibility of equally adequate models of an evolutionary event. On the surface, this view is unsurprising. Scientific models provide partial, or idealized, descriptions of reality. They underscore certain features while obscuring others. Thus, constructing a model entails choosing which causal features to make salient at the expense of others. The existence of multiple accurate representations of an evolutionary event is an implication of the pluralism with respect to causality endorsed above.[7] If there are multiple, true causal stories involved at different levels of the biological hierarchy in an

[7] By "accurate model," I mean that the model truthfully represents the phenomenon at hand, and not a model that produces 'accurate' data. For instance, the Geocentric Model of the cosmos yielded accurate data, but that data did not truthfully align with the universe.

evolutionary situation, then it makes sense that these causes can be modeled at different levels as well. This is the version of pluralism advanced in this paper—different models of an evolutionary event are available to us *because* the complicated causal stories can be told at different levels.[8]

Importantly, the existence of alternative correct models is not problematic so long as the representations in question are not incompatible, meaning that they do not conflict with empirical matters of fact. This is true despite that different models will lead to different implications. For instance, ascribing fitness to a cell might yield different insights than ascribing fitness to a genome. This plurality of models should not be taken as a fatal limitation of science; rather, it should be explored to its logical extreme to uncover what insights different models can provide—how does treating the genome, or the cell, or the multicellular unit as an interactor in a selection process underscore causally relevant details that other models overlook? Having elaborated on the two aspects of the pluralism I am defending here relating to causality and alternative models, my task in the remainder of this paper is to vindicate this view by showing how different models of cancer evolution can first, truthfully represent the selection process in cancer, and second, yield practical implications for researchers.

[8] This view takes the ability to accurately model an evolutionary event at a plurality of levels to be sufficient grounds to claim that the interactor at one of those levels is a true target of selection. In other words, if cancer can be truthfully modeled as genomic selection, then the genome can be reasonably considered a target of selection or the level at which selection acts. Robert Brandon (1982, 1988, 1990) has argued that an additional condition is required for there to be selection at a particular level, namely, that the phenotypes at that level must "screen off properties of entities at every other level from reproductive values at a given level" (1990, pg. 88). I address Brandon's analysis in section 6.2.

5. How Different Models Can Truthfully Represent Cancer Evolution

There are two aspects to understanding how different models can truthfully represent an evolutionary process. The first, to use Helen Longino's (2006) wording, involves parsing the causes differently. Pluralists believe there are multiple ways to tell the causal story, and that nature cannot be divided at unique joints. As Waters (2005) writes, "scientists and philosophers who believe the proper aim of science is to discover the comprehensive theory that provides the single, correct way to represent the causal structure of the world (or part of the world) need to temper their realism." The second aspect concerns where different models draw the conceptual divide between the environment and the selected entity. For instance, genic selectionists such as Williams (1966) and Dawkins (1982) draw this divide between an allele and everything outside of the allele, while group selectionists such as Sober and Wilson (1998) draw the divide between groups of organisms and everything beyond the groups. As a result, because causes can be parsed in different ways, models that draw the environmental divide at different levels of the biological hierarchy can each isolate real and different causal stories. With this in mind, I now turn to showing how a genomic and multicellular model can each truthfully represent cancer evolution. Because I take the cellular model outlined in section 1 to be self-evident, I do not discuss it here.

5.1 *A Genomic Model*

A genomic model draws the environmental divide between the genome and everything outside of it, including the genetic environment contained within a cell. Hence, instead of treating the cancer cell as the selected entity, this model considers the cancer genome as the level at which selection is occurring. Suppose there are two different genomes in a tumor: $genome_1$ and

genome$_2$. W_1 and W_2 represent the selection coefficients for each genome, which are measures of the relative fitness of the genomes. The relative values of the selection coefficients are $W_1 > W_2$. To explain why genome$_1$ is selected for, one must take into account the causal processes that lead to an increase in its relative frequency. Say genome$_1$ possesses a mutation that genome$_2$ does not have, and this alteration confers a selective advantage to genome$_1$ by allowing it to induce angiogenesis, leading to an increase in its frequency. Angiogenesis, the formation of new blood vessels that acquire extra nutrients and oxygen for cellular processes, appears to consist of causal processes that occur mainly at the cellular level. A genomic model does not deny the importance of these processes. In fact, a genomic model must take these important cellular processes into account in order to accurately represent the evolutionary process. However, it does not need to model angiogenesis as interactions between a selected entity and its environment. Rather, a genomic model treats the cellular causal processes of angiogenesis as part of the environment external to the genome. It adjusts where the environment is drawn.

Of course, in order to represent selection at the genomic level, there must be a causal relationship between the increase in relative frequency of genome$_1$ and the heritable variation between genome$_1$ and genome$_2$. This is no problem at all for the pluralist. One only needs to show that genome$_1$ has properties that causally affect its ability to leave copies of itself. To show this, however, I must first introduce a theory of causation that can identify the difference between genuine causes and pseudocauses.

James Woodward's manipulability theory of causation offers the conceptual tools to accomplish this task. On Woodward's account:

> The claim that *X* causes *Y* means that for at least some individuals, there is a possible manipulation of some value of *X* that they possess, which, given other appropriate conditions (perhaps including manipulations that fix other variables distinct from *X* at certain values), will change the value of *Y* or the probability distribution of *Y* for those individuals (2003, pg. 40).

The main feature of Woodward's theory is that causal claims involve counterfactual dependencies that allow us to answer 'what if' questions regarding what would happen if certain properties were manipulated. While a detailed discussion of the intricacies of Woodward's account is beyond the scope of this paper, it will be helpful to clarify a special type of manipulation that Woodward calls an 'intervention.'

An intervention is "an ideal experimental manipulation of the value of *X* performed for the purpose of determining whether *X* causes *Y*" (2003, pg. 14). Thus, an intervention involves manipulating one single variable while keeping all other variables constant. This allows us to rule out misleading cases in which a manipulation of the value of *X* that changes the value of *Y* does not constitute a causal relation. For instance, consider the statement: if the barometer reading were to fall, a storm would occur. An intervention here would involve some process that alters the value of the barometer reading in a way causally and probabilistically independent of all other variables. If this intervention led to a storm, then under Woodward's account the manipulation of the barometer reading caused the storm. Of course we know that this intervention would not lead to a storm, and hence there is no causal relationship between altering the barometer reading and the occurrence of a storm. The converse is not true. An intervention that altered the weather conditions to resemble a storm would lead to the barometer reading falling, so it is appropriate to say that the storm caused the barometer reading to fall. For our

purposes, it will be enough to remember that *X* causes *Y* if an intervention in the value of *X* would change the value of *Y*.

We can now return to the issue of how different genomes have properties that causally affect their ability to leave copies of themselves. Recall that $genome_1$ has a mutation in a gene that induces angiogenesis. Because this model considers the genome as the selected entity, the trait that affects fitness would be the mutated sequence itself. Suppose there is a genetically homogeneous small tumor that is early in the multistage development of cancer. At this point all of the cells are genetically identical and in the same microenvironment. This sets up a helpful model to see how intervening on the genome could yield different effects in a causal manner. Say a mutation occurs in the VEGF-A gene on one of the genomes. The VEGF-A gene encodes ligands that help orchestrate new blood vessel growth (Ferrara, 2009). The mutation leads to the constitutive overexpression of VEGF-A, resulting in an increase in blood vessel growth in the cellular environment and thus an increased ability of the genome to leave copies of itself. After the mutation, we have two genomes and the only difference between them is the mutation in $genome_1$. Under Woodward's account, then, the mutation would be a *cause* of the increase in frequency of the genome, because an intervention that altered the genome and nothing else increased the fitness of $genome_1$. And, treating the mutation as a trait of the genome, we have differential reproductive success of the genome that is a causal result of the variation in traits.[9]

It is worth noting that this model of selection is representing the same evolutionary event that a cellular model would. It just does so by considering traits to be at the genomic level. The basic thesis is that "the causes of one and the same selection process can be correctly described

[9] Monists might object that Woodward's theory of causation is mistaken. However, it is impossible to deny that there is still some causal relation between the mutation in $genome_1$ that leads to angiogenesis and an increase in the reproductive success of $genome_1$. If Woodward's theory does not accurately capture this causality, then the burden falls on the monists to clarify what it means to be a 'cause' in a way that does not undermine their thesis.

at different levels" (Waters, 1991). As we can see under Woodward's manipulationist theory of causation, intervening at the genome level can get the causal story right. On this model, we are treating the genome as an interactor in the evolutionary process.

5.2 *A Multicellular Model*

A multicellular model draws the environmental divide between the multicellular unit and everything outside of it. This model is complicated because it is essentially a model of group selection, which has historically been controversial in the evolutionary biology literature. Group selection is the claim that natural selection can act on groups with properties that are not explainable by the properties of their individual components, and therefore the existence of groups irreducible to the existence of their components must be conceded. The primary catalyst for models of group selection has been the need to account for altruism. In evolutionary biology, altruism is defined as a behavior that costs the actor and benefits the recipient of that action in terms of expected fitness. Although the necessity to explain the presence of altruism in cancer biology poses a less urgent problem since the progression of cancer is well documented unlike the history of life, group models of selection are still both explanatorily and practically valuable for representing cancer evolution. In this section, I consider two documented instances of cancer that suggest the appropriateness of group selection models. The first sheds light on the possibility of a group adaptation in cancer, and the second concerns the structure of tumor subpopulations**.** The literature on group selection is vast, with different proponents adopting different criterion for what constitutes group selection. My aim here is not to assess the validity of this historically controversial approach to modeling selection. Rather, my aim is to first, point to what appear to be group-like characteristics in some cases of cancer evolution, and second, to suggest that these

characteristics are sufficient grounds to accept a model of multicellular selection given its practical implications, which I show in the last section.

Recent studies have suggested the existence of a group adaptation that causally increases the fitness of a multicellular unit. Aceto *et al.* (2014) found that cells that escape from tumors sometimes do so in clusters, and that these clusters have a higher chance of forming lung metastases when compared to single cells that escaped. These clusters of blood-borne circulating tumor cells (CTCs) possess a metastatic potential 23 – 50 times that of single CTCs. In addition, when comparing the gene expression differences between single and cluster CTCs, Aceto *et al.* found no large differences other than the expression of cluster-associated genes such as plakoglobin, which is a protein that plays a pivotal role in regulating cell-cell adhesion and possesses the ability to affect cell signaling (Aktary and Pasdar, 2012). Interestingly, when plakoglobin expression in tumor cells was reduced, the CTC clusters dissociated, resulting in a decrease in metastatic ability (Figure 3). *Prima facie*, we have something that strongly resembles group selection here—if groups of cells are better at invading and surviving in distant tissues than single cells, then it appears that the 'group' is selected for in this case.

The mechanisms mediating plakoglobin adhesion activity in metastasis are still unclear, but it seems plausible to assume that there is something advantageous about being a multicellular unit that enables it to better colonize distant organs than a single CTC. Is this a group adaptation? I think so, because the collective fitness of the multicellular unit is higher than the average of the individual fitness of the particles alone. To see why this is true, suppose that a cluster of CTCs, and the individual cells genetically identical to those cells comprising that cluster, were injected into the bloodstream of mice. If the cluster has a higher fitness than the mean fitness of the individual cells, then there seems to be some group-level property allowing the cluster to have a

higher relative fitness than the average of the single CTCs. Indeed, when Aceto *et al.* injected both single cells and clustered cells into the tail vein of mice, they found that the clustered cells were more resistant to apoptosis, so it appears that the average fitness of the constituent cells will not be able to account for the collective fitness. This finding seems to beckon for a group selection model.

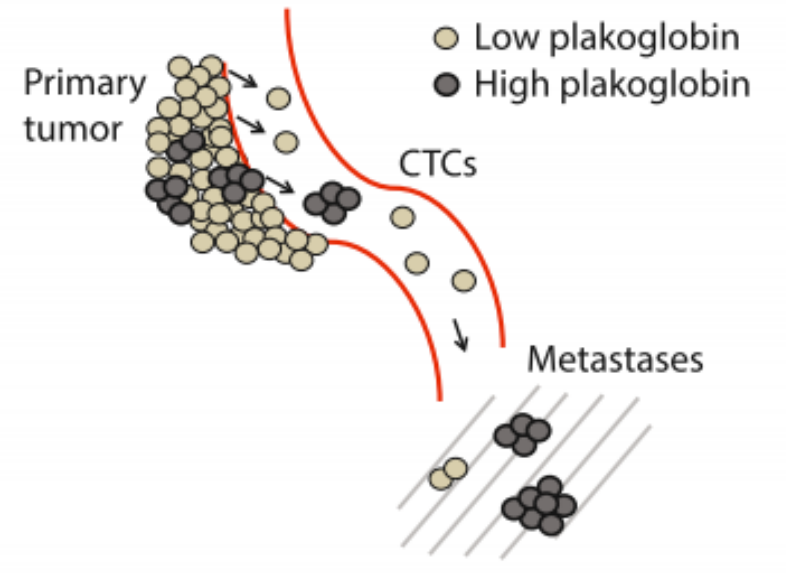

Figure 3. Cells in the primary tumor that express high plakoglobin are likely to generate CTC clusters with elevated metastatic potential. [Reprinted from Aceto *et al.* 2014.]

Another reason that group selection models for cancer may be fitting is the structure of intratumor heterogeneity in cancer progression. Importantly, this heterogeneity is not so much manifested in the individual cells themselves so much as in *groups* of cells. Tumors are composed of genetically distinct subpopulations of cells that compete against each other for resources. Hence, considering a multicellular unit as the fitness bearing entity in modeling cancer evolution is quite appropriate given the subclonal divisions within a tumor. These competing subclones are significant to cancer progression for a number of reasons. First, overall tumor

growth can sometimes be driven by a single minor cell subpopulation. Marusyk *et al.* (2014) found that certain subclones, initially present as a minority competing against parental cells within a tumor, could enhance expansion of both the subclone and the parental cells. In addition, they found that tumors with subclones present at an initial 1:18 ratio competing against each other grew faster than monoclonal tumors. This suggests that the presence of biological interactions between different subpopulations can lead to new cancer phenotypes such as increased proliferation. Second, these interactions between subpopulations of cells can be understood, from an evolutionary perspective, as instances of cooperation or cheating.[10] While intratumor heterogeneity can be maintained by selection for the fittest subclone, these subclones can also cooperate by sharing growth factors or by depending on each other for survival. However, this mutual interdependence, as might be expected, is not always the case. Archetti *et al.* (2014) found that cells that do not produce growth factors can free-ride on the growth factors secreted by nearby producer cells. This "tragedy of the commons," as it is called in game theory, leads to a delicate group dynamic analogous to the equilibrium in cooperative societies such as ant colonies—if the number of cheaters becomes too high, then the cooperators cannot successfully fulfill the tasks for the society to sustain itself. These recent studies suggest two instances in which it might be appropriate to ascribe fitness to multicellular units, rather than individual cells. Specifically, a tumor can be modeled as an environment within which there is intergroup competition between the heterogeneous subclones, or the tumor itself can be considered a group with a delicate equilibrium that sustains its survival.

One might argue that group models are unnecessary since a cellular account can provide all the relevant explanatory and causal details. That is, one can parse the causes in a way such

[10] Cooperation here is not necessarily altruistic. Two cells may mutually benefit from cooperating without either of them suffering a cost in fitness.

that the cell is the unit of selection, and the group or other cells surrounding it are simply part of the environment. On this account, each cell would be ascribed a different fitness that takes into account any group-level interactions as part of the environment. This argument against group selection models is valid, and not new. However, the validity of this critique does not contradict the pluralism I am advancing here. On the contrary, it actually supports it, for the presence of two different perspectives that are not empirically contradictory would vindicate pluralism. Indeed, Kerr and Godfrey Smith (2002) have shown that individual and group level perspectives in a multilevel selection scenario can be mathematically equivalent.[11] It is important to note that I have not argued for group selection models to the exclusion of other models. This would be monism, which I reject. Rather, I have argued for the possibility that group selection models can accurately capture the relevant dynamics in certain cases of cancer evolution, and thus should be further explored in conjunction with other available models.

[11] Kerr and Godfrey-Smith acknowledge that their conclusion says nothing about causality, for one might argue that in any given scenario, only one of the two parameterizations is the correct representation of the underlying causal facts, even though they are mathematically equivalent. Nevertheless, the work of Kerr and Godfrey-Smith does suggest that pluralism is at the very least a coherent possibility.

6. Implications of Pluralism for Cancer Research

Pluralism about the levels of selection raises a second question: on what grounds do we choose between competing models? Waters (2005) claims these grounds are pragmatic. He writes:

> Scientific models provide partial descriptions. They highlight some features and obscure others. In simple situations, models may highlight all causally relevant features and obscure the rest. In complicated situations, however, such as biological development and evolution, models typically make some causal features salient by obscuring others. This doesn't render models false, but depending upon scientists' interests, it makes them more or less suitable, pragmatically speaking.

While I think Waters is correct to claim the choice between different models can be made on practical grounds, he does not expand on how this might be carried out. In this section, I aim to extend Waters' argument by outlining the pragmatic advantages of cellular, genomic, and multicellular models in studying the evolution of cancer. The biggest pay off of pluralism is the ability to adopt different perspectives on an evolutionary process. In a field such as cancer research, taking into account the benefits of all available models is a currently overlooked tool in the scientific literature that should be taken more seriously.

6.1 *Implications of a Cellular Model*

The cellular model of cancer is certainly the most intuitive. After all, it is the cell that is actually cancerous; it is the cell that is interacting with the microenvironment; it is the cell that develops resistance to drug therapy. Virtually every current therapy—chemotherapy, radiation therapy, targeted cancer therapies, angiogenesis inhibitors, biological therapies—is premised on decreasing the fitness of the cancer cell ("Types of Treatment," National Cancer Institute), and framing cancer progression as a cellular process has yielded many therapeutic insights.

Perhaps the two most important insights include studying resistance to treatment and the application of ecological theory to cancer. Cancer therapies positively select for cells that possess resistance to treatments such as chemotherapy. If this resistance is heritable, then as those without resistance are killed, the environmental niche clears for the resistant cells to proliferate. This positive selection is the underlying cause of recurrence. Quantitative genetics tells us that the rate of evolution of a trait is proportional to the selective differential (Falconer and Mackay, 1996). When chemotherapy fails to destroy all of the cancerous cells, the difference in the value of the resistance trait between the surviving cells and the killed cells increases. It follows, then, that therapies with high (but not complete) cancer cell mortality increase the evolutionary rate of resistance. Consequently, attempts to eradicate cancer often fail to be effective long-term treatments. There have been various evolutionary approaches to avoid this problem, such as selecting for benign cells, altering the carrying capacity of the neoplasm, and provoking competition between neoplastic and normal cells (Merlo *et al.*, 2006). These avenues still need to be further explored, but they exemplify how focusing on the evolutionary causes of cancer at the cellular level may offer superior strategies in treating cancer.

The application of ecological theory to cancer therapy has also yielded new avenues to explore in the treatment of cancer at the cellular level. For example, Gatenby (2009) suggests that our strategy should not be focused on finding a 'cure' for cancer, but rather should be aimed at finding a way to control disease. Much like how pest management aims not to eradicate a species, but rather to restrict the population growth, cancer treatments should be designed not to destroy the cancerous cells, but to maintain a stable, tolerable tumor (Figure 3). The reason for this lies in the emergent resistance to treatment, which dramatically reduces a patient's chances of survival. Gatenby's suggestion exemplifies precisely how an understanding of the cancer cell

as an interactor in a selection process is invaluable to the present, and future, of cancer research. Applying this principle requires researchers to look beyond the proximate causes, and instead focus on understanding the mechanisms by which cancer cells acquire resistance, how these mechanisms can be exploited, and how the evolutionary dynamics of the tumor microenvironment can be altered to achieve this end.

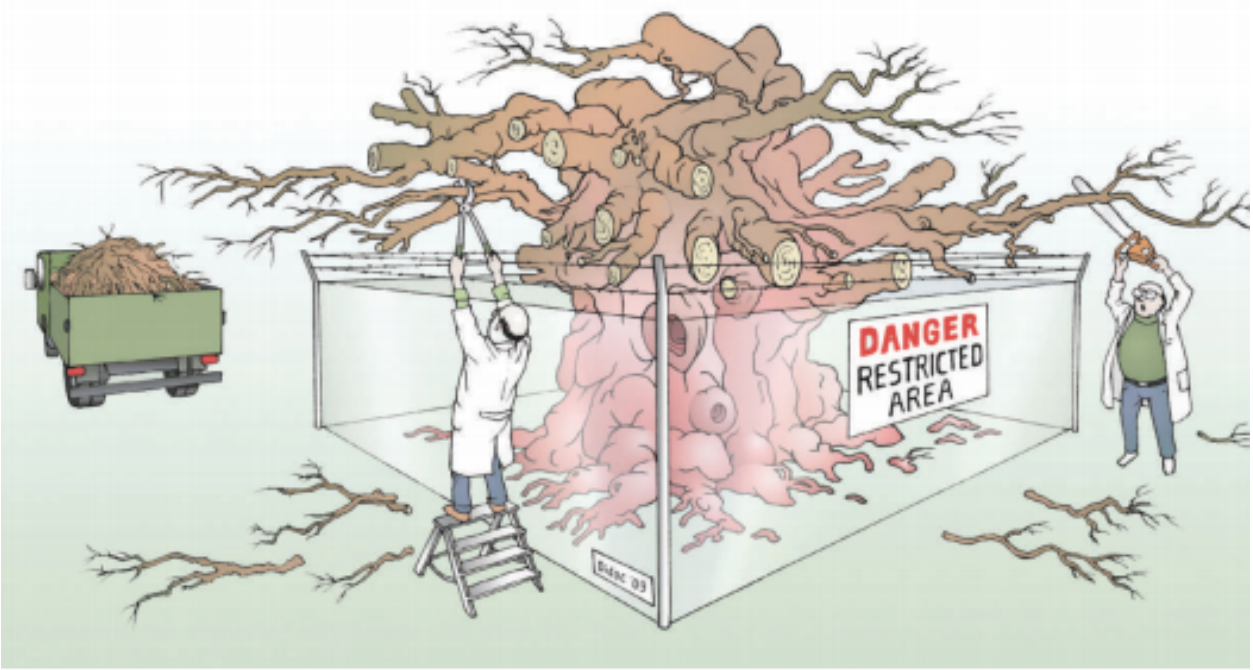

Figure 4. Illustration of the strategy Gatenby (2009) proposes. Cancer maintenance may be a better therapeutic approach than cancer eradication because the cost of resistance to treatment dramatically decreases the chances of patient survival. [Reprinted from Gatenby, 2009.]

6.2 *Implications of a Genomic Model*

While the current model of cancer as cellular selection is invaluable, it lumps together the diverse sets of genes that give rise to cancer into convergent phenotypes at the cellular level characterized by the Hallmarks of Cancer. This model treats two cancer cells that exhibit an equal degree of sustained proliferative signaling resulting from mutations in different sets of genes between the two cells as equally fit, all else being equal. However, this would be

obscuring important causal features at the genetic level. One set of genes might also confer resistance to chemotherapy, while the other may not. A cellular model would not be able to differentiate this genetic diversity until the advent of clinical intervention. The benefit of a genomic model is that it identifies these deeper genetic causes of resistance that are invisible to the cellular model. And it does so without obscuring what the cellular model highlights. The genomic model simply extends the concept of the environment further down the hierarchy and draws the division at the genome and everything outside of it. This does not obscure the important processes that are occurring in the genomic environment such as angiogenesis. It just parses the causes differently. Hence, the valuable insights from ecological theory that entail altering the environment to affect the fitness of the selected entity are still relevant and applicable (Gatenby, 2009). The genomic model just focuses on the fitness of the genome instead of the cell.

The lumping obstacle I mentioned above that the cellular model faces is particularly relevant in the clinic. Cancer is characterized by genetic heterogeneity across different cancers as well as within individual cancers. To date, nearly 500 cancer genes have been identified (Futreal *et al.*, 2004). In some cancers, specific genes are consistently mutated. For instance, retinoblastoma is characterized by inactivating mutations in retinoblastoma 1 (*RB1*), a tumor suppressor gene (Blanquet *et al.*, 1995). However, genetic signatures of cancers are exceptions rather than rules. The majority of common cancers are associated with a number of different cancer genes that are mutated at a low frequency. For example, a recent genome-wide analysis of 489 high-grade serous ovarian cancers identified thousands of somatic mutations, yet only 10 of those were recurrently mutated cancer genes (Ellis *et al,* 2012). There is, philosophically speaking, multiple realizability of the Hallmarks of Cancer. Of particular significance, genetic

heterogeneity manifested within a single tumor has clinical implications for how evolution proceeds—the more genetic variation within the population, the more likely it is to possess mutants with selective advantages that can withstand therapeutic intervention. A genomic model would underscore this genetic variability by treating each distinct genome as a separate entity, competing for reproductive success within the microenvironment.

A genomic model also provides a unified framework that can bridge molecular biology and evolutionary cancer biology. Molecular biologists interested in uncovering, for instance, the genetic regulation of cell proliferation cannot accommodate a cellular model of cancer evolution into their work because their focus is on the mechanisms internal to the cell itself. Treating the genome as an interactor in an evolutionary process can avoid this issue, as well as provide a more intuitive and explanatorily superior model of understanding cancer evolution for molecular biologists.

To see why this is true, consider a counter argument to the genomic model, and more broadly, to pluralism about the levels of selection. Robert Brandon (1982, 1988, 1990) has argued that the concept of screening off can be used to provide the conditions required for selection to act at a given level. The basic idea of screening off is this:

> If A renders B statistically irrelevant with respect to outcome E but not vice versa, then A is a better causal explainer of E than is B. In symbols, A screens off B from E if and only if $P(E,A \cdot B) = P(E,A) \neq P(E,B)$ [read '$P(E,A \cdot B)$' as the probability of E given A and B]. If A screens off B from E then in the presence of A, B is statistically irrelevant to E, that is, $P(E,A) = P(E,A \cdot B)$.
> (Brandon, 1990, pg. 83)

Brandon uses screening off to illustrate the point that natural selection does not favor phenotypes and genotypes equally. Rather, phenotypes but not genotypes are directly visible to natural selection, since phenotypes screen off both genes and genotypes when explaining the reproductive success of an interactor. Note that if Brandon's use of screening off to answer the

levels of selection question is correct, then there is little to no room for a pluralistic thesis. Screening off delivers the conceptual tools to identify the 'real' level at which selection acts in a hierarchy of entities.

Brandon's starting point is that for there to be selection at a particular level, it is not enough for entities at that level to have heritable and differential reproduction as a result of variances in fitness. What is needed, in addition, is that the differences in fitness have the right sort of 'causal explainer,' namely, that 'the "phenotypes" of the entities at that level screen off properties of entities at every other level from reproductive values at the given level" (Brandon, 1990, pg. 88). A pluralist would accept the first criterion as sufficient grounds for there to be selection at a particular level, but reject the second criterion relating to phenotypes screening off properties of entities at every other level. Brandon's analysis entails that the genomic model presented in this paper is not a true level at which selection occurs. This is because the cellular phenotypes screen off the underlying genetic mutations. To use our earlier example in section 5.1, the cellular phenotype of increased blood vessel growth, or angiogenesis, screens off the mutation in the VEGF-A gene, and thus selection is actually occurring at the cellular level since angiogenesis is the phenotype on which selection is acting.

One way to respond to Brandon's analysis is to point out that on a genomic model, the phenotypes are considered at the genomic level already, since we are treating a mutation as a trait of the genome. Thus, the phenotypes do not screen off the genotypes because the interactor and replicator on this model are the same entity. But this response is arguably circular, for it presupposes the phenotypes to be at the genomic level in order to avoid the reasonable alternative that the phenotypes are actually at the cellular level. A second and more powerful response is to argue that even 'if *A* renders *B* statistically irrelevant with respect to outcome *E*,'

that is, even if *A* screens off *B*, it does not necessarily follow that '*A* is a better *causal explainer* of *E* than is *B*' (emphasis added). Cancer provides a particularly relevant example to illustrate why sometimes a genomic level phenotype is actually a superior causal explainer to a cellular level phenotype.

Consider the clinical value of taking advantage of synthetic lethality in the design of targeted therapies. Synthetic lethality is when mutations in two or more genes lead to cell death, yet a mutation in only one of these genes does not. Of particular significance, the sensitivity of $BRCA1^{-/-}$ cells to poly(ADP-ribose) polymerase (PARP) inhibitors is a well-known synthetic lethality that is utilized in cancer therapy (Farmer *et al.*, 2005).[12] *BRCA1* is necessary for homologous recombination repair of double strand DNA breaks. Although $BRCA1^{-/-}$ cells can still survive despite this genetic defect, the cells must depend on alternative repair pathways that involve PARP function. Hence, in $BRCA1^{-/-}$ cells that are targeted with PARP inhibitors, DNA breaks cannot be repaired, leading to cell arrest and death. Nevertheless, often times during or after PARP inhibitor therapy, 'reversion' mutations in *BRCA1* can partially restore the functionality of the *BRCA1* protein complex, allowing the mutant clone to avoid cell death (Swisher *et al.*, 2008).

In the scenario above, we can understand and explain the occurrence of resistance to PARP inhibitors solely at the genomic level and without reference to any cellular phenotypes. Not only can the explanans be given solely in molecular terms, but it is also *explanatorily superior* to any cellular level explanans because it identifies all of the causally relevant facts, namely, the deeper genetic causes of resistance to PARP inhibitor therapy. A cellular model would focus on higher-level causes that obscure these lower-level details. The view that genetic

[12] $^{-/-}$ is the symbol used to denote that both copies of a gene are recessive, or in this case, defective.

factors are privileged causal explainers is, I think, a reasonable one. And philosophers of biology have given good reasons to support this idea. Rosenberg (2006), for instance, advocates for a gene-centric research program grounded in molecular biology that can avoid the untenable dualism between physicalism and antireductionism.

In the above example, it is also more intuitive to consider the genome as an interactor because the ability or inability to restore DNA strand breaks is a trait that is better modeled at the genomic level than the cellular level from a pragmatic standpoint for molecular biologists. Consider how the genome interacts with the environment here. The PARP inhibitors constitute an environmental factor that decreases the fitness of the genome by affecting other pathways in the *genetic* environment. This interaction between the genome and environment accommodates the focus of molecular biologists to exploit the complex genetic networks implicated in cancer for therapeutic purposes, parsing the extracellular factors that are highlighted in a cellular model as more distant causes. For the molecular biologist, this genomic model is preferable because it is more consistent with their research focus—they are not so much interested in the ecological interactions between different cells so much as the ways those ecological dynamics alter the underlying molecular mechanisms in the genetic environment. Thus, even though Brandon is right to claim that a cellular phenotype would screen off a genomic level one, this does not entail that the cellular phenotype is a better causal explainer, at least in cancer evolution. A genomic model can sometimes highlight all of the causally relevant details without reference to any higher-level facts, and offer a causal story that is explanatorily superior to a cellular level account. A genomic model also provides a more intuitive representation of the evolutionary situation relative to the aims of a molecular biologist.

In the near future, the goal of research will be to elucidate how the cancer genome evolves and responds to treatment. A patient's cancer will be sequenced and characterized during diagnosis. The implicated cancer genes and mutations will be identified and compared to a large-scale database of past genomic analyses. As these databases containing clinical outcome data linked to genomic analyses reach a sufficiently large sample size, oncologists will be able to develop algorithms that can help choose a personalized therapeutic approach. Modeling cancer evolution as evolution of the genome will be instrumental to this end.

6.3 *Implications of a Multicellular Model*

In section 5.2, I argued for the applicability of multicellular models of selection in cancer evolution by pointing to recent studies—Aceto *et al.* (2014) reported the presence of a group adaptation in clusters of CTCs, Marusyk *et al.* (2014) found that interactions between competing groups of cells in a tumor can lead to new tumor phenotypes such as increased proliferation, and Archetti *et al.* (2014) reported the presence of cancer cells that free-ride on the growth factors secreted by neighboring cells. In this section, I elaborate on the implications of these findings.

The findings of Aceto *et al.* (2014) are an important step in identifying the underlying causes of metastasis. This is especially significant given that the metastatic spread of breast cancer, usually to the bone, lungs, liver, or brain, accounts for the large majority of cancer-related deaths (Nguyen *et al.*, 2009). Thus, their report that CTC clusters have a 23-50 fold increased metastatic potential compared to single CTCs paves the path to a novel therapeutic target, namely, the clusters themselves. Traditionally, studies of cancer metastasis have focused on the concept of "seed versus soil" as an important determinant of metastatic potential (Fidler, 2003). This model focuses on the match between the successful metastatic cell (the 'seed') with

its mutated signaling pathways that confer proliferative and invasive properties, and a specific microenvironment (the 'soil') that can facilitate metastatic growth. While the "seed versus soil" concept is undoubtedly important, it largely focuses on the genes and proteins that are associated with metastasis. Aceto *et al.* (2014) offer a different perspective on metastasis by underscoring the importance of the physical traits that benefit a cluster in the bloodstream. They write, "the physical characteristics of single CTCs and CTC clusters may also contribute to metastatic propensity, especially as they impact the ability of epithelial tumor cells to survive the loss of cell adherence and shear forces in the blood stream" (pg. 1110). These two approaches to understanding metastasis—the "seed versus soil" concept and the emphasis on physical characteristics of clusters—suggest two alternative ways to clinically disrupt CTC clusters. On the "seed versus soil" model, one possibility is to disrupt the match between the 'seed' and the 'soil' by targeting specific cluster-associated genes, such as plakoglobin. Since plakoglobin knockdown abrogates intercellular interactions and inhibits the ability to form clusters, as Aceto *et al.* (2014) demonstrate, this particular genetic pathway can be interrupted in therapy to reduce the ability of the cluster to invade foreign tissues. Alternatively, a therapeutic intervention that aims to physically interfere with the stability of clusters in the bloodstream can also be further explored. This approach might take advantage of recent technological advancements for CTC capture to screen for clusters based on physical properties such as size, density, and electromechanical characteristics in order to remove or dissociate the clusters (Yu *et al.*, 2011). These two approaches represent different ways to decrease the fitness of the multicellular unit.

The findings of Marusyk *et al.* (2014) also suggest the value of using a multicellular model. Their report emphasized the relevance of cross-talk between groups of cells within a tumor. They write:

> Our results suggest that tumors can be driven by a sub-population of cells that does not have higher fitness, but instead stimulate growth of all tumor cells non-cell-autonomously by inducing tumor-promoting microenvironmental changes. Conversely, non-cell-autonomous clonal expansion does not necessarily translate into increased tumor growth rates. The non-cell-autonomous driver subclone can be outcompeted by a sub-clone with higher proliferative output, thus collapsing the tumor (Marusyk, *et al*, 2014).

These results highlight the importance of the evolutionary dynamics between subclones within a tumor population, and suggest that clinical diagnostics that focus only on the most abundant subpopulation of cells may be misguided. Their findings point to two possible therapeutic approaches. One approach would be to identify and alter the 'tumor-promoting microenvironmental changes' brought about by a specific subclone that is driving tumor growth. Another approach would be to increase the fitness of a subclone and allow it to outcompete the driver subclone, leading to tumor collapse. Both of these approaches will need to take into account not only the interactions between groups of cells within a tumor, but also the interaction within a particular group.

Finally, the results of Archetti *et al.* (2014) also emphasize the value of a multicellular model, but in a slightly different way. They documented that pancreatic cancer cells that do not produce insulin-like growth factor II (IGF-II) are at a proliferative disadvantage in pure cultures (where exogenous IFG-II is not present), but at a proliferative advantage in mixed cultures (where exogenous IGF-II is present), because they can free-ride on the IFG-II secreted by producer cells. Their work suggests the value of considering the tumor itself as a group, or a society with the dynamics of cooperation and defection. This model is important therapeutically

because cooperation between tumor cells remains one of the largest obstacles for treatments that target growth factors. Of particular importance, therapies that reduce the amount of circulating growth factors may only lead to a temporary reduction in tumor growth. This is because a decrease in circulating growth factors leads to an increased number of IFG-II producer cells necessary to achieve a benefit for the overall tumor population. Consequently, the equilibrium shifts to a higher fraction of producer cells, which may explain the relapse observed in treatments that target growth factors (Bergers and Hanahan, 2008).

Nevertheless, treating a tumor as a society of cooperating and defecting cells suggests at least two possible therapeutic insights. First, the dynamics of the production of growth factors can be modified by, for instance, increasing the diffusion range. As Archetti *et al.* (2014) show in their model, the increased diffusion range leads to an increased size of the group that benefits from the growth factors, resulting in an increased cost per benefit ratio for the producer cells. With a higher cost of production, producer cells decrease in fitness, which reduces the fraction of them in a tumor and hence the growth rate of the tumor itself. Second, the dynamics of the society can be altered to let the cheating cells win by outcompeting the producer cells in the tumor. Archetti *et al.* (2014) found that in populations that began with a high proportion of cheaters and low proportion of producers, the cheaters would grow so quickly that eventually the producers would go extinct. This points to a new therapeutic strategy that entails, in effect, creating a tumor within a tumor. That is, by taking out the producer cells from a tumor, genetically modifying them to no longer secrete growth factors, and then reinserting these newly created cheaters back into the tumor, scientists may be able to disrupt the tumor dynamics so badly that the tumor collapses. This approach is a spin-off of the "tragedy of the commons" in game theory—since growth factors are an intratumor public good, the cheater's ability to free-

ride on other members of the group's contributions leads to a fitness advantage that will eventually result in the demise of the population due to the lack of the public good.

7. Conclusion

In this paper I argued for pluralism by proposing that genomic, cellular, and multicellular models of cancer can all get the causal story of cancer evolution right. This was because the causal chains involved are so complicated that they can be cut up in different ways. I also suggested that these models each have pragmatic advantages by being able to distinguish important causal factors that other models overlook or cannot account for. Hence, each model paves a different avenue to the development of novel therapies. Further, since cancer is an extraordinarily dynamic disease, one of the benefits of a pluralistic approach to studying the evolution of cancer is that it can capture the diversity of ways in which different cancer cell types undergo carcinogenesis and metastasis. If in fact there are a plurality of models that accurately capture the causal chains in cancer evolution, then the big payoff, as Kitcher, Sterelny, and Waters (1990) put it, is that "philosophers and biologists can turn their attention to more serious projects than that of quibbling about the real unit of selection." Cancer researchers do not need to take sides on what the 'real' unit of selection is, and limit themselves to a cellular model of cancer evolution as has been done in the scientific literature. It is enough to adopt useful heuristics that can maximize the outcome of therapy. In this work I have tried to show how this might be carried out, and more broadly, how philosophical insights can have important consequences for the practice of science.

References

Aceto, N. *et al.* (2014). Circulating tumor cell clusters are oligoclonal precursors of breast cancer metastasis. *Cell.* 158(5): 1110-1122.

Aktary, Z. and Manijeh, P. (2012). Plakoglobin: Role in Tumorigenesis and Metastasis. *International Journal of Cell Biology.* 2012.

Aktipis, C.A., Kwan, V., Johnson, K.A., Neuberg, S.L., Maley, C.C. (2011). Overlooking Evolution: A Systematic Analysis of Cancer Relapse and Therapeutic Resistance Research. *PLoS ONE.* 6(11): e26100.

Aktipis, C.A. and Nesse, R. (2012). Evolutionary foundations for cancer biology. *Evolutionary Applications.* 6: 144 – 159.

Archetti, M., Ferraro, D.A., Christofori, G. (2014). Heterogeneity for IGF-II production maintained by public goods dynamics in neuroendocrine pancreatic cancer. *PNAS.* 112(6): 1833- 1838.

Bergers, G. and Hanahan, D. (2008). Modes of resistance to anti-angiogenic therapy. *Nat Rev Cancer*. 8(8): 592 – 603.

Bignell, G.R. *et al.* (2010). Signatures of mutation and selection in the cancer genome. *Nature*. 463: 893-898.

Blanquet, V. *et al.* (1995). Spectrum of germline mutations in the RB1 gene: a study of 232 patients with hereditary and non hereditary retinoblastoma. *Hum. Mol. Genet*. 4(3): 383-388.

Brandon, Robert. (1978). Adaptation and evolutionary theory. *Studies in History and Philosophy of Science*. 9(3): 181 – 206.

Brandon, Robert. (1982). The Levels of Selection. *PSA.* 1: 315-323.

Brandon, Robert. (1988). The Levels of Selection: A Hierarchy of Interactors. In H.C. Plotkin (ed.). *The Role of Behavior in Evolution*. Cambridge, MA: MIT Press, 51 – 71.

Brandon, Robert. (1990). *Adaptation and Environment.* Princeton: Princeton University Press.

Cairns, J. (1975). Mutation selection and the natural history of cancer. *Nature*. *255*:197–200.

Caulin, A.F., and C.C. Maley. (2011). Peto's Paradox: evolution's prescription for cancer prevention. *Trends in ecology & evolution.* 26: 175 – 182.

Cohen, D. (1985). The canine transmissible venereal tumor: a unique result of tumor progression. *Adv Cancer Res.* 43: 75-112.

Dawkins, R. (1978). Replicator Selection and the Extended Phenotype. *Z. Tierpsychologie.* 47: 61 – 76.

Dawkins, R. (1982). *The extended phenotype: the gene as the unit of selection*. Oxford: Oxford University Press.

Dugatkin, L.A. and Reeve, H.K. (1994). Behavioural Ecology and Levels of Selection: Dissolving the Group Selection Controversy. *Advances in the Study of Behavior.* 23: 101-33.

Ellis, M.J. *et al.* (2012). Whole-genome analysis informs breast cancer response to aromatase inhibition. *Nature.* 486: 353-360.

Falconer, D.S., and T.F.C. Mackay. (1996). *Introduction to Quantitative Genetics.* Longmans Green, Harlow, Essex, UK.

Farmer, H, *et al.* (2005). Targeting the DNA repair defect in *BRCA* mutant cells as a therapeutic strategy. *Nature.* 434: 917 – 921.

Fidler, I.J. (2003). The pathogenesis of cancer metastasis: the 'seed and soil' hypothesis revisited. *Nat Rev Cancer.* 3: 453 – 458.

Franklin-Hall, L.R. Explaining Causal Selection with Explanatory Causal Economy: Biology and Beyond. In P.A. Braillard & C. Malaterre (eds.). *Explanation in Biology: An Enquiry into the Diversity of Explanatory Patterns in the Life Sciences.* Springer (forthcoming).

Futreal, P.A *et al*. (2004). A census of human cancer genes. *Nature Rev. Cancer.* 4: 177-183.

Gatenby, R. (2009). A change of strategy in the war on cancer. *Nature*. 459: 508-509.

Greaves, M. (2012). Cancer stem cells as 'units of selection.' *Evolutionary Applications.* 6: 102 – 108.

Greaves, M. and Maley, C. (2012). Clonal evolution in cancer. *Nature*. 481: 306-313.

Greenman, C., Wooster, R., Futreal, P.A., Stratton, M.R. and Easton, D.F. (2006). Statistical analysis of pathogenicity of somatic mutations in cancer. *Genetics.* 173: 2187 – 2198.

Griesemer, J. (2000). The Units of Evolutionary Transition. *Selection*. 1: 67-80.

Hanahan, D. and Weinberg, R.A. (2011). Hallmarks of cancer: the next generation. *Cell. 144(5)*: 646-74.

Hawkins CE, Baars C, Hesterman H, GJ H, Jones ME, Lazenby B *et al*. (2006). Emerging disease and population decline of an island endemic, the Tasmanian devil. *Sarcophilus harrisii*. *Biol Conserv.* 131: 307–324.

Kerr, B. and Godfrey-Smith, P. (2002). Individualist and Multi-level Perspectives on Selection in Structured Populations. *Biology and Philosophy.*17: 477-517.

Kim, R., Emi, M., and Tanabe, K. (2007). Cancer immunoediting from immune surveillance to immune escape. *Immunology.* 121: 1-14.

Kitcher, P. (2004). Interview with Philip Kitcher. *Human Nature Review 4.* 87-92. <http://human-nature.com/nibbs/04/kitcher.html>

Kitcher, Philip, Kim Sterelny, and C. Kenneth Waters. (1990). The Illusory Riches of Sober's Monism. *Journal of Philosophy.* 87: 158 – 160.

Lachish S, Jones M, McCallum H. (2007). The impact of disease on the survival and population growth rate of the Tasmanian devil. *J Anim Ecol.* 76: 926–936.

Lloyd, E. (2001). "Units and Levels of Selection: An Anatomy of the Units of Selection Debates", in R.S. Singh, C. B. Krimbas, D. B. Paul, and J. Beatty (eds.) *Thinking About Evolution* – Vol. 2 – Historical, Philosophical, and Political Perspectives. Cambridge University Press.

Lloyd, E. (2005). Why the Gene Will Not Return. *Philosophy of Science.* 72(2): 287 – 310.

Loh R, Bergfeld J, Hayes D, O'Hara A, Pyecroft S, Raidal S *et al*. (2006a). The pathology of devil facial tumor disease (DFTD) in Tasmanian Devils (*Sarcophilus harrisii*). *Vet Pathol.* 43: 890–895.

Longino, H.E. (2006). Theoretical pluralism and the Scientific Study of Behavior in *Scientific Pluralism.* Eds. Kellert, S., Longino, H., and C. K. Waters. Minneapolis, MN: University of Minnesota Press.

Marusyk, A., Tabassum, D.P., Philipp, M.A., Almendro, V., Franziska, M., Polyak, K. (2014). Non-cell-autonomous driving of tumour growth supports sub-clonal heterogeneity. *Nature.* 514: 54-58.

Maynard Smith, J. and Szathmáry, E. (1995). *The Major Transition in Evolution*. Oxford: Oxford University Press.

McCallum H, Tompkins DM, Jones M, Lachish S, Marvanek S, Lazenby B *et al*. (2007). Distribution and impacts of Tasmanian devil facial tumor disease. *EcoHealth.* 4: 318–325.

McCallum H. (2008). Tasmanian devil facial tumour disease: lessons for conservation biology. *Trends Ecol Evol (Personal edition).* 23: 631–637.

McShea, D.W. (1996). Metazoan Complexity and Evolution: is there a Trend? *Evolution.* 50: 477-492.

McShea, D.W. (1998). Possible Large-Scale Trends in Organismal Evolution: Eight Live Hypotheses. *Annual Review of Ecology and Systematics.* 29: 293 – 318.

McShea, D.W. (2001). The Hierarchical Structure of Organisms: a Scale and Documentation of a Trend in the Maximum. *Paleobiology.* 27: 405 – 423.

Merlo, L., Pepper, J., Reid, B., Maley, C. (2006). Cancer as an evolutionary and ecological process. *Nature Reviews Cancer*. 6: 924-935.

Michod, R.E. (1999). *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality.* Princeton, NJ: Princeton University Press.

Mill, J. S. (1846) *A System of Logic*. New York: Harper & Brothers.

Murchison, EP. (2009). Clonally transmissible cancers in dogs and Tasmanian devils. *Oncogene.* 27, S19-S30.

Nguyen, D.X., Bos, P.D., and Massagué, J. (2009). Metastasis: from dissemination to organ-specific colonization. *Nat Rev Cancer.* 9: 274 – 284.

Nowell, PC. (1976). The clonal evolution of tumor cell populations. *Science*. *194*: 23–28.

Okasha, S. (2006). *Evolution and the Levels of Selection*: Oxford University Press, USA.

Rosenberg, A. (2006). *Darwinian Reductionism: Or, How to Stop Worrying and Love Molecular Biology.* Chicago: University of Chicago Press.

Sober, E. and Wilson, D.S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.

Sober, E. and Wilson, D.S. (2002). Perspectives and Parameterizations: Commentary on Benjamin Kerr and Peter Godfrey-Smith's "Individualist and Multi-level Perspectives on Selection in Structured Populations". *Biology and Philosophy.* 17(4): 529-37.

Sterelny, Kim and Philip Kitcher. (1988). The Return of the Gene. *Journal of Philosophy.* 85: 339-361.

Strevens, M. (2008). *Depth.* Cambridge, MA: Harvard University Press.

Swisher, E.M., *et al.* (2008). Secondary *BRCA1* mutations in *BRCA1*-mutated ovarian carcinomas with platinum resistance. *Cancer Res.* 68: 2581 – 2586.

Teng, M.W.L., Swann, J.B., Koebel, C.M., Schreiber, R.D., and Smyth, M.J. (2008). Immune-mediated dormancy: an equilibrium with cancer. *J. Leukoc. Biol.* 84: 988-993.

Types of Treatment. National Cancer Institute. Retrieved April 18, 2014, from http://www.cancer.gov/cancertopics/treatment/types-of-treatment

Waters, C.K. (1991). Tempered Realism about the Force of Selection. *Philosophy of Science.* 58: 553 – 573.

Waters, C. K. (2005). Why Genic and Multilevel Selection Theories Are Here to Stay. *Philosophy of Science.* 72: 311 – 333.

Waters, C. K. (2007). Causes that make a difference. *Journal of Philosophy.* 104: 551 – 79.

Williams, George. (1966). *Adaptation and Natural Selection.* Princeton, NJ: Princeton University Press.

Yates, L.R. and Campbell, P.J. (2012). Evolution of the cancer genome. *Nat Rev Genet.* 13(11): 795 – 806.

Youn, A. and Simon, R. (2011). Identifying cancer driver genes in tumor genome sequencing studies. *Bioinformatics.* 27: 175 – 181.

Yu, M., Stott, S., Toner, M., Maheswaran, S., and Haber, D.A. (2011). Circulating tumor cells: approaches to isolation and characterization. *J. Cell Biol.* 192: 373 – 382.