

Is There a Right Not to Know Genetic Information about Oneself?

Benjamin E. Berkman

Abstract

The move from targeted genetic testing to genomic sequencing has produced a number of ethical debates, but the most controversial question is the extent to which individuals have a right not to know genetic information about themselves. This chapter explores the extent to which it is ethically necessary to respect someone's choice to remain deliberately ignorant about this kind of information. Challenging the majority view that there is a nearly absolute right not to know, arguments are presented which push back against that vigorously held (although not always rigorously defended) position, in support of the idea that we should abandon the notion of a strong right not to know. Drawing on the fields of bioethics, philosophy, and social science, an extended argument is provided in support of a default for returning high-value genetic information without asking about a preference not to know. Recommendations are offered about how best to balance individual autonomy and professional beneficence to guide the field of genomic medicine as it continues to evolve.

Introduction

The field of bioethics is replete with cases where people choose to remain deliberately ignorant about information. In the clinical realm, patients regularly refuse to learn their diagnosis or prognosis, thus making it difficult to engage them in making informed decisions about their medical care. Some patients ignore efficacy and side-effect information when deciding which medication to take. People choose not to read consent forms before enrolling in research. To avoid bias in research data collection, investigators agree to be left in the dark about whether their subjects have received an active intervention or the inert placebo. But nowhere in bioethics has the concept of deliberate ignorance proven to be more controversial than in the genetic testing sphere. Using this

as a case study, I will illustrate in this chapter some of the fascinating tensions inherent in discussions about deliberate ignorance, and provide a real-world example to help think through the normative and policy implications when people actively choose not to learn information

With its promise to revolutionize medicine, the capacity to cheaply and quickly generate an individual's entire genome has prompted a series of significant ethical debates. Genomic sequencing, unlike targeted genetic testing, produces massive amounts of extraneous information, some of which can have relevance for an individual's health. The mismatch between the specific indication that led to the ordering of the test and the breadth of results that the test produces has ignited an ongoing debate about the ethics of managing *incidental* or *secondary findings* (Presidential Commission for the Study of Bioethical Issues 2013); that is, information (typically clinically significant and medically actionable) that is generated during a test or procedure but which does not relate to the original purpose for which the test or procedure was conducted (Wolf et al. 2008).

Over the last decade, the problem of managing incidental or secondary findings has been a major source of contention in the research ethics and science policy realms. The most contentious part of the debate has focused on the extent to which it is ethically necessary to respect a person's *right not to know* (RNTK) genetic information about themselves. To make this problem more concrete, imagine the following scenario: As part of a diagnostic workup for what is suspected to be a rare genetic disorder, a patient undergoes genome sequencing. Prior to the procedure, during the informed consent process, the patient clearly checks the box indicating the choice to opt out of receiving incidental genetic results. When the physicians analyze the resulting genomic data, they find evidence of a high genetic risk for a different disorder, hereditary nonpolyposis colon cancer (HNPCC). Since HNPCC is treatable if found early, but is nearly always fatal if discovered at late stages, they recognize the intrinsic value of this information to the patient. It could enable the patient to seek enhanced screening for a cancer that is very difficult to detect with normal colonoscopies and, in turn, prevent serious disease and even save the patient's life. Should the physicians disclose this finding to the patient, despite the explicit choice made by the patient not to be informed of secondary findings?

Whether or not a patient's RNTK needs must be honored under all conditions has sparked a highly contentious debate. The conflict highlights a classic problem in bioethics: the frequent tension between autonomy and beneficence. Our society places an extremely high value on empowering and honoring an individual's choices, particularly in the medical realm. This often presents a clear dilemma for physicians, who want to act in a way that provides the highest prospect of benefit for their patients. In this case, it means asking whether a patient's choice not to know should be honored at the cost of an opportunity to take advantage of potentially beneficial medical information.

As genomic sequencing technology continues to be fine-tuned and implemented, the examination of ethical norms and standards of care requires serious deliberation. Is the RNTK appropriate in a genomic era, given the obvious and inevitable conflict between autonomy and beneficence that such a right creates? Because the ability to control what genetic information is revealed has been imbued with the power of a right, debate has thus far been unduly focused on the seemingly absolute nature of an individual's autonomy. The majority view among bioethics scholars seems to be that the RNTK continues to be of paramount importance and should not be abrogated in any way.

A case can be made, however, that genomic medicine is adhering too tightly to an outdated conception of the RNTK. My goal in this chapter is to push back against that vigorously held (although not always rigorously defended) position, in defense of the idea that the notion of a *strong* RNTK should be abandoned. I will offer an extended argument in support of a default for returning high value (defined below) genetic information *without asking* about a preference not to know.

Emergence of the Controversy

Researchers and bioethicists have been grappling with the problem of genetic incidental findings for over a decade. From this debate, which has been both protracted and often quite heated, the RNTK emerged as an uncontroversial issue, at least initially. As commentators argued about the circumstances under which there was an obligation to return individual findings, and which findings to return, there seemed to be broad support for the view that findings should only be returned when the research participant expressed the desire to receive this information (Fabsitz et al. 2010). In terms of an "obligation" on the part of researchers, that obligation was to *offer* individual findings to research subjects, who could elect to receive or refuse the information. Accordingly, there was wide agreement that researchers should discuss the RNTK with potential subjects and prospectively solicit their binding preferences.

These early views on the RNTK were expressed in the nascent days of genomic medicine, before large-scale genomic sequencing emerged. As sequencing technology advanced, and particularly as it moved from the research setting into the clinical realm, the debate began to shift for two related reasons. First, the utility of genomic sequencing was improving. An increasing number of genetic variants had been strongly linked to a range of phenotypes where knowledge of one's genetic status could have a profound impact on treatments for (or prevention of) serious disease. Second, a growing number of patients were being sequenced, leading to reasonable projections about the important role that genomic sequencing would have as a regular part of clinical care.

In response, the American College of Medical Genetics and Genomics (ACMG) issued recommendations for the reporting of incidental findings in clinical exome and genome sequencing (Green et al. 2013). Their goal was to start a conversation about clinical standards for managing the predictable onslaught of medically relevant incidental findings. These recommendations suggested that labs should actively search (i.e., opportunistically screen) for a “minimum list” of variants that predispose patients to risk for disorders that “would likely have medical benefit for the patients and families of patients undergoing clinical sequencing.” Considering both the weight of the scientific evidence and the clinical implications of knowing the genetic information, the ACMG limited the list to “unequivocally pathogenic mutations in genes where pathogenic variants lead to disease with very high probability and where evidence strongly supports the benefits of early intervention.” Controversially, the ACMG Working Group argued against soliciting patient preferences about receiving (or not receiving) incidental findings. They did not think that it was appropriate to give patients a choice not to learn about clinically important and actionable findings, advancing the claim that clinicians have a fiduciary duty to warn patients about high-risk variants where an intervention is available. Ironically, this argument against a patient’s strong RNTK actually involves clinicians actively blinding themselves to patient preferences.

The recommendation against soliciting patient preferences for not knowing genetic information ignited an extended (and often quite spirited) debate within the research ethics community. A relatively small set of commentators tried to defend the call for mandatory disclosure of high-value incidental findings (Berkman and Hull 2014; McGuire et al. 2013). The overwhelming majority view, however, was extremely critical of the recommendation, holding that patients have a strong RNTK and that any abrogation of that right was inappropriate (Burke et al. 2013; Wolf et al. 2013). As Trinidad et al. (2015) stated, the ACMG statement was “an instance of paternalistic overreach” that should be “widely rejected as inconsistent with the ethical and legal duties of clinicians.” Even more interesting was the fact that these pro-RNTK views were often couched in absolute terms. Commentators were not blind to the fact that strongly preferencing the RNTK meant that some patients might not receive information that could save their lives. Although not expressed in exclusively principistic language, these arguments essentially seem to advance the view that autonomy should override beneficence in RNTK situations.

In response to the mounting criticism of their recommendations, ACMG retreated from their initial position. Citing a purported consensus among ACMG members, the organization refined their position to state that before a sample is sent for analysis, patients should be allowed to opt out of receiving incidental findings. I believe that this majority view is mistaken and argue, in the remainder of this chapter, that there should not be a strong RNTK high-value genetic information about oneself.

What Is High-Value Genetic Information?

The first step in my analysis is to define the type of information that I will subsequently argue should not be subject to the RNTK. To be clear, I do not intend to argue that there is no role for patient preferences in determining when to receive any genetic information. Rather I focus on the extent to which there is a RNTK (a) *high-value* genetic information where (b) medical action can mitigate or prevent mortality or serious morbidity when (c) there is strong evidence of the link between genotype and significant disease risk. The arguments made in this chapter should not be directly applied to information where there is no medical action to take (e.g., Huntington disease), when the condition is less severe (e.g., asthma), or when the evidence is weak (e.g., single case reports). In essence, I will be arguing against the RNTK with the relatively small set of potential findings on the ACMG list in mind.

To illustrate how valuable information from the ACMG list might be, it is useful to consider some of the variables laid out by Schwartz et al. (this volume). They enumerate a number of features that can make a given piece of information more or less valuable. In constructing its list, the ACMG intentionally selected only variants associated with serious diseases, so the *magnitude* of the information's importance would necessarily be high. Schwartz et al. defined their list to include only *actionable* conditions, where there is an effective preventative or medical intervention to take. They also mitigated uncertainty by choosing only those variants that had a *high-quality evidence* base and a high penetrance meaning, such that a given finding in a particular individual would likely be *decisively* relevant to that person given their age and clinical presentation.

It is also important to consider the possible harms associated with revealing this information. Critics have expressed concern about a number of possible risks. Most prominent were psychosocial concerns such as stigma, discrimination, and anxiety (Klitzman et al. 2013; Lázaro-Muñoz et al. 2015). There were also worries about the iatrogenic and economic impact of unnecessary follow-up procedures and interventions, both on individual and population levels (Burke et al. 2013; Klitzman et al. 2013; Wolf et al. 2013). This second set of concerns was predicated on the prior predictive value problem; since existing evidence is based on studies involving affected families, critics argued that it is premature to assume similar penetrance in families without a history of the disease because there could be as yet unidentified mitigating genetic features that could reduce or eliminate risk (Holtzman 2013). Invoking the precautionary principle, these critics argued that we should avoid returning incidental information until we can be more certain that doing so will not prompt unnecessary medical interventions. As I explain below, there are reasons to think that these concerns are not as significant as commentators assume and that they are outweighed by the potential benefits of knowing the information.

The Philosophical Origins of the RNTK: An Unexpectedly Contested Concept

Before analyzing the benefits and risks, it is useful to begin by considering the philosophical origins of the RNTK, as this reveals the concept to be on much more analytically shaky ground than many contemporary commentators acknowledge. The RNTK genetic information is a relatively new idea: it first appeared in the literature in the 1970s and 1980s but did not really gain traction until the 1990s (Laurie 1999, 2000; Takala 1999). A substantial body of work developed in the subsequent decade, concurrent with the gradual incorporation of genetic testing into clinical medicine. While there appears to be significant recent support for the RNTK, a robust examination of the concept must begin with an analysis of the idea's philosophical origins. Contemporary RNTK advocates have tended to present their views in the absence of this historical perspective, seemingly arguing that a strong, autonomy-based RNTK is self-evident (Herring and Foster 2012). In contrast to this assumption, I believe that a close examination of the earlier RNTK literature reveals a much more controverted and nuanced history. Specifically, I will demonstrate that acceptance of a strict RNTK is far from universal in the philosophical literature, and that even staunch proponents recognize that the RNTK can be easily overridden by competing considerations.

Arguments for a Strong RNTK

Most commonly, scholars ground the RNTK in autonomy, arguing that one's right to self-determination implies a right to make decisions about learning (or not learning) sensitive medical information. These scholars typically build their argument on a foundational assertion that genetic information has the potential to cause psychological and economic harm (Andorno 2004). While often granting that more information can allow for improved decision making regarding future plans, they stress that for some individuals, this information can lead to anxiety, depression, stigma, and possible discrimination. Therefore, an individual should be afforded the freedom to weigh the risk of psychosocial and economic harms against the potential benefit that the knowledge might provide.

Beyond this basic argument, RNTK proponents often cite concerns about paternalism in medical care, making claims that it has no place in modern medicine, even if justified by seemingly reasonable considerations (Takala 2001). They often draw a distinction between preventing harm and creating benefit, arguing that knowing one's genetic status offers the possibility of a benefit, but not knowing about a genetic defect does not directly harm the person, since the defect is present regardless. If forced provision of information poses a risk of harm and there is only a possibility of creating benefit (rather than prevention of harm), unwanted provision of genetic information is indefensibly

paternalistic. Accordingly, it is paternalistic to overrule individual choice, even if a choice differs from our conception of what is “reasonable.”

Interestingly, there does not seem to be overwhelming support in the foundational RNTK literature for a strong, autonomy-based RNTK. The limited number of scholars cited above support such a view, but the weight of the literature is squarely against an expansive view of the RNTK. As I will discuss in the next few sections, most scholars either argued for a much narrower conception of the RNTK or dismissed the idea entirely.

Arguments against a Strong RNTK

Autonomy Misapplied

A number of arguments challenge the notion of a strong, autonomy-based RNTK. One main strain of criticism asserts that the concept of an autonomy-based RNTK is too broad and that the principle has been misapplied. As Rhodes (1988:433) argued, “misunderstandings about the nature and moral force of autonomy have led some in the genetics community to a false conclusion about genetic ignorance.” The concern with an autonomy-based RNTK stems from the commonly held view that rights are “preemptive and value-laden,” and thus the content of a right must be carefully articulated and defended (Laurie 2014).

Autonomy fails as a basis for doing so for a number of reasons. First, a strong, autonomy-based RNTK is inappropriate because that line of reasoning unrealistically requires ignoring the fact that there is no such thing as an unfettered choice (Harris and Keywood 2001). There are lots of things that people would like to do (or not do) or to know (or not know), but one sometimes must make less than ideal choices. Second, there is no basis for the idea that information alone is autonomy-constraining, because a clear distinction can be drawn between obtaining relevant information and making subsequent decisions on the basis of that information. Finally, autonomy is not boundless; there are certain actions that are prohibited as a matter of public policy.

Instead of autonomy, Laurie (2014) builds his RNTK theory on the idea of spatial privacy, or the notion that we have a right to ensure that an individual is “in a state of non-access.” Spatial privacy includes both the familiar notion of physical separateness and encompasses separateness of the individual’s psyche. The latter form of privacy entitles an individual to protect his or her own sense of self. As a result, it can be an invasion of one’s “psychological spatial privacy” to receive information about oneself that was not already possessed.

Ultimately, in basing his conception of a RNTK in privacy rather than autonomy, Laurie calls into question the view that the RNTK is actually a strong right. Laurie acknowledges that a decision to violate someone’s psychological privacy involves a number of competing factors and must be holistically

assessed instead of being held up as a strict ethical rule. Even if unwanted disclosure constitutes a violation of privacy, it can be justifiable under certain circumstances. Ultimately, he (and others) argue for a *prima facie* presumption in favor of the RNTK. This presumption can be rebutted, however, in a range of clinical cases. When considering the justifiability of violating someone's psychological integrity, a number of factors should be relevant, including the availability of a cure or intervention, severity of the condition, and the likelihood of disease manifestation.

The Incoherence Objection

In addition to challenging the notion that the principle of autonomy plausibly supports a RNTK, some critics make an even more forceful argument, calling into question the very coherence of the RNTK as a concept. These scholars make an autonomy claim of their own, but in the opposite direction, advancing the idea that knowledge is necessary in order to exercise autonomy (Malpas 2005). One needs to know that there is an issue that requires a decision; so not knowing undermines one's ability to make an autonomous choice. Rather, autonomy demands "critical reflection," which includes thoughtful, informed decision making and deliberation. Without relevant information, it is impossible to make informed decisions about future plans, and ill-informed decisions may even frustrate one's future self, as an individual may make choices that are ultimately self-defeating.

Some have gone as far as saying that autonomy requires rationality and freedom of will, but patients who deny themselves readily available information are not acting rationally, as they are depriving themselves of relevant health information (Ost 1984). If someone is so fixed in their intentions that no amount of relevant information would change their mind, this would be tantamount to an irrational obsession. Similarly, it is logically impossible for someone to claim to know *a priori* that information will not be relevant to his or her decision. This line of reasoning not only rejects a RNTK, it seems also to imply a moral duty to be informed about information that would make a difference in decisions (at least when it can be obtained without undue effort).

Effects on Third Parties

A third objection to a strict, autonomy-based RNTK is founded on a concern about the effect on others of maintaining one's ignorance. According to this line of reasoning, genetic information unavoidably involves relatives, and one has an obligation to learn readily available information about one's health to enable relatives the opportunity to act on that knowledge. Relatives who had not previously known about a familial genetic risk would also be able to benefit from knowing by taking a variety of actions, such as seeking their own genetic

testing, changing risk-associated behaviors, pursuing prophylactic treatment options, and engaging in rigorous screening (Bottis 2000).

An individual's desire to refuse genetic information may conflict with the duty to warn a potentially affected family member. This does not mean that the RNTK should not be respected if possible, but rather that there are clear situations where other competing ethical principles might cause one to disregard a desired RNTK. On this view, the RNTK should not be viewed as a strict right; when there is a conflict between the RNTK and the right of relatives to sensitive genetic information concerning their own health, the RNTK must yield, due to the very real risk of harm to the family members.

Outdated Examples

Finally, defenders of the RNTK often point to concerns about testing for Huntington disease (HD) and Alzheimer disease (AD), citing data on people's reluctance to get tested to illustrate the potential anxiety that people feel when faced with negative genetic information and to support the claim that there is a very real risk of harm associated with unwanted provision of one's genetic status (Austad 1996). Both examples, however, stem from the targeted genetics era, and I would argue that they have limited utility as valid comparators in the modern genomic era. HD and AD are devastating and presently inimitable neurological conditions. As such they are *sui generis* since they present the possibility of psychological harm without any corresponding clinical benefit. When these kinds of examples are utilized by scholars, they should only be used to make a claim about the RNTK genetic information associated with commensurate diseases. But this isn't the case; commentators consistently use these limited examples to make broader claims. Citing evidence of concern about being tested for HD or AD is irrelevant to this important debate, since the real empirical and normative questions relate to whether people would or should refuse to learn about potentially life-saving genetic information.

Moving Away from a Strong RNTK

In the previous section, I took a close look at the philosophical origins of the RNTK. Contrary to what contemporary commentators have been arguing, the notion of a strong, autonomy-based RNTK rests on an unstable conceptual foundation. Only a handful of philosophers have endorsed such a position, with the majority either arguing for a much more limited, nonautonomous-based conception, or even against the whole concept entirely. Here, I borrow from the bioethics and social science fields to make additional arguments for abandoning the notion of a strong RNTK.

I begin by reframing the debate away from an autonomy-dominated perspective, providing a comprehensive analysis of the harms and benefits that

result from adhering to a strong RNTK position. From this analysis, I conclude that the potential health benefits of abandoning a strong RNTK greatly outweigh the concomitant harms, thereby challenging the idea that psychosocial concerns should automatically get to trump the prospect of life-saving intervention. I end by exploring two additional considerations that are relevant to any rigorous discussion of the RNTK: moral distress and genetic exceptionalism.

Analyzing the Impact of a Strong RNTK

There is reason to believe that people's views on the RNTK are less settled than one might have previously believed; while autonomy and the RNTK may seem sacrosanct in isolation, forcing people to confront the trade-offs inherent in real-world scenarios changes many minds (Gliwa et al. 2016). If people are open to considering trade-offs between autonomy and beneficence, then it becomes important to rigorously examine what those trade-offs might entail. This kind of analysis has thus far been absent from the RNTK debate. The overwhelming focus in the recent literature on an autonomy-based RNTK has had the unfortunate effect of short-circuiting discussion of the topic by directing attention solely on the harms associated with not honoring individual preferences.

One can see this in the arguments in favor of a strong RNTK, which generally focus on patient autonomy, appealing to the long history of shared medical decision making and respect for patient preferences. For example, as an impressively credentialed group of bioethicists forcefully argued (Burke et al. 2013:857):

...choice matters. Patients may wish to decline the additional analysis on a number of grounds....Concepts of shared decision making and respect for patient preferences argue for offering meaningful choices wherever possible, with appropriate information to allow patients to choose the best option for themselves.... If patients decline additional testing, it follows that the laboratory should not perform the additional analyses.

In addition, critics often supported their strong autonomy arguments by claiming that there is good reason to think that many people do not want to learn certain kinds of genetic information about themselves (Jarvik et al. 2014; Klitzman et al. 2013). Furthermore, commentators were comfortable with the idea that people should even be able to refuse information with profound medical significance. For example, as Wolf et al. (2013:1050) put it:

Patients have the right to refuse testing and findings, even if potentially lifesaving. Just because many patients might want this information does not mean that it can or should be imposed on all.

Similarly, a number of commentators cite the legal right to refuse medical interventions, arguing that an individual's ability to place limits on treatments also implies a legal right to refuse medical information.

This laser focus on autonomy has not allowed for a comprehensive analysis of the harms *and* benefits of honoring *or* ignoring the RNTK. The reality is that any policy will have potential negative consequences. Whichever option is chosen, we will necessarily be making a mistake in one of two directions: unwanted disclosure or lost opportunity for medical intervention. Here, I lay out what I take to be the full set of relevant considerations and explore some of the relevant empirical data that can help us fully assess the overall impact of any RNTK policy. Specifically, there are three empirical questions that should be carefully considered, which I explore below.

How many people genuinely do not want to know genetic information about themselves, if it could have a profound impact on morbidity or mortality?

Available data support the reasonable claim that the overwhelming majority of people would want to be given genetic risk information that will have a direct impact on their health. In one study, nearly all respondents wanted to learn about a range of genetic risk factors, with 90% wanting to learn about nonactionable health risks and 96% wanting to learn about actionable genetic risk factors (Kaufman et al. 2008). Similarly, in the largest study to date of views toward the return of incidental findings resulting from sequencing research, nearly 5,000 members of the public were surveyed and nearly all of them (98%) wanted to learn about genetic risk for life-threatening conditions that can be prevented (Middleton et al. 2016). A strong majority even wanted to know about life-threatening conditions that could not be treated. So as a baseline, it seems fair to say that the vast majority of people would actively want to know high-value health information, although more research is needed to establish the real-world contours of this claim since these data are based on surveys that asked people to respond to theoretical scenarios.

Of course, that leaves a very small subset of the population who might not want to know this information. Although this is an empirical question that requires further study, it is plausible that this small set of people who would not want to know is primarily comprised of individuals for whom clinical action might not be indicated (e.g., patients with a terminal illness, the elderly, people with a religious objection to receiving medical treatment). Proponents of the RNTK point to these types of examples in defense of their views. My counterargument is that these relatively rare examples should not drive the RNTK debate; we should not be creating a broad RNTK policy based on a limited set of cases where the medical information actually has little or no value to the individual. Rather, these cases can be addressed separately, because they represent scenarios where doctors can reasonably anticipate a need to actively solicit preferences. Doctors should be able to predict most cases where an individual

patient might have good reason to not know information because it is not clinically actionable for them given their situation. Even if there are some cases where doctors cannot easily predict that a patient has a reason for not wanting to know, if that reason is strong enough, those patients will likely self-identify.

If people were given genetic risk information that they would have preferred not to know, what is the magnitude of the harm they actually experience?

Psychological Harm. If the vast majority of people would want to know important genetic risk information, and if most of those who would not want to know can be bracketed, we are arguably left with an exceedingly small set of people. More empirical research is needed to ascertain the exact size and composition of this group, but whatever that number turns out to be, the next task is to examine the magnitude of harm that this small group will experience if given information that they would have preferred not to know. As discussed above, RNTK proponents frequently make claims about the danger of psychological harms flowing from the disclosure of negative genetic information. These claims rely on limited data related to a few poorly targeted examples, such as HD and AD. What can the broader psychological literature tell us about our reactions to unfortunate genetic information?

The short answer is that psychological research has demonstrated that people are not as good as they think at affective forecasting or predicting the magnitude and duration of our future emotional reaction to both positive and negative events (Wilson and Gilbert 2005). For example, recent lottery winners typically overestimate the length and duration of their spike in happiness. Similarly, but in the opposite direction, people who have recently lost a loved one overestimate the length and duration of their negative emotional response to the traumatic event. In both cases, after an initial spike, people gradually tend to return to their previous baseline level of happiness. Essentially, the mind is assumed to have a sort of psychological immune system, which helps people handle negative information, often making the actual impact of negative information significantly smaller than the expected negative impact. However, when making a prediction about future emotional responses, we disregard our ability to cope, thereby overestimating the negative impact of information.

This is particularly true in the medical realm, where the literature suggests that an individual's predictions concerning the emotional consequence of learning about genetic disease risk do not square with people's actual ability to adapt to negative health information (Halpern and Arnold 2008). In a broad range of medical contexts, there are data showing that the affective forecasting bias is particularly pronounced when healthy people are asked to assess the negative emotional impact of (theoretical) future health problems. Specifically, people generally assume that receiving negative genetic information will be devastating, but research demonstrates that people are much better at coping with negative information than they think they will be. In reality, we should

be careful about assuming that the negative psychological effect of receiving risk information for many untreatable conditions is as significant as many assume. More studies are needed to ascertain how a broad range of people react to negative news, but the existing evidence suggests that we should be open to the idea that negative reactions to unfortunate genetic information will be relatively mild and transient (Broadstock et al. 2000).

It is striking that RNTK proponents continue to make claims about the harmful psychological impact of genetic information when there is such limited empirical support for such concerns. More evidence about emotional reactions to genetic information would certainly be useful, but the existing literature at least raises important questions about whether we “systematically overestimate the durability and intensity of the affective impact of events on well-being,” thereby creating a “culture of risk-aversion in which patients may be opting out of potentially beneficial diagnostic and treatment regimes” (Peters et al. 2014:312).

Economic Harm. If psychological harm seems less likely and serious than is often assumed, there is still the issue of economic harm (i.e., discrimination). The likelihood and magnitude of discrimination is somewhat more difficult to assess, but existing data suggests that perhaps there is less cause for concern than previously thought (Rothstein 2008). It does appear that there are occasional instances of discrimination in these realms, but that they are primarily associated with untreatable single gene conditions (e.g., HD) that carry little weight for purposes of determining whether there should be a broad RNTK. Even with some scattered evidence of discrimination in these realms, a systematic review of existing data calls into question the need for a policy intervention (Joly et al. 2013), suggesting that there is a significant gap between the fears of genetic discrimination and actual reality.

Again, this is not to suggest that genetic discrimination will never become a problem in life or in terms of long-term care insurance. Rather, my argument is that we should make a clear-eyed assessment of the frequency and magnitude of any economic harms flowing from disclosure of genetic risk information before automatically assuming a worst-case scenario. As I explore in the next section, there are some potential negative effects associated with honoring a strong RNTK, which should be balanced against a rigorous evaluation of the harms associated with not doing so.

What Is the Cost of Always Soliciting Patient Preferences?

On one side of the scale, we have a very small group of people who are arguably at very low risk of experiencing significant, lasting psychological or economic harm. On the other side, we would want to know the impact of adopting a robust RNTK policy that involved actively soliciting individual preferences. My argument is that such a policy would necessarily result in some loss of opportunity to provide people with valuable information because there is good

reason to doubt our ability to assess, accurately and reliably, people's true preferences.

A number of arguments support this claim. The first concerns how people engage with informed consent documents. Extensive data suggest that people frequently do not carefully read consent forms, and when they do, that their understanding and appreciation of the content can often be lacking (Mandava et al. 2012). If subjects are signing consent forms with such incomplete understanding of the important details contained therein, it seems questionable to have confidence in the infallibility of any process designed to solicit preferences about knowing genetic incidental findings. This is particularly true given the inherent complexity of genetic information and the associated difficulty patients will have in making a choice in that context. Many commentators have expressed concern that the wide range of types of genomic findings will be overwhelming and could become a significant barrier to implementing truly informed consent (McGuire and Beskow 2010).

In the pre-genomic era when targeted genetic testing was the norm, patients could reasonably absorb the range of information they might receive; a single gene test typically only revealed information associated with the relevant condition. Now, when genome sequencing is employed, it is impossible to know what kind of results will be generated, making the informed consent process that much more difficult. Ensuring patient comprehension and managing expectations becomes increasingly difficult as the amount of genomic data generated grows. Furthermore, it will even be difficult to adequately describe the variety of genomic information categories because of terminological confusion. Terms such as "actionability," "clinical utility," and "clinical significance" are typically used to describe the types of findings someone might or might not desire, but there is a lack of conceptual clarity about exactly what those terms mean (Eckstein et al. 2014).

There are also concerns about how preferences can shift over time. Life events and the passage of time can change a person's views; an answer given as a single young adult might differ to one that the same person would give once they are married with children. Unless the medical world can develop a process for actively re-soliciting preferences (an unrealistic proposition), there is the very real risk that a binding decision made at a single point in time could become inconsistent with future desires.

Informed consent is a cornerstone of bioethics, and with good reason. In its ideal form, it allows doctors and researchers to demonstrate respect for persons and allows competent individuals to make autonomous choices about their engagement with medicine. The arguments made above should not be read as a wholesale indictment of informed consent. Rather, my point is that we should be skeptical about our capacity to assess, adequately and accurately, individual preferences about knowing or not knowing specific categories of genetic information. There is a very real risk that a policy of actively soliciting preferences about knowing or not knowing genetic information could result

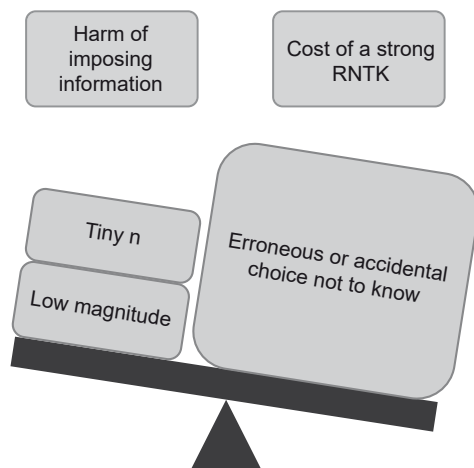


Figure 12.1 A framework for comprehensively analyzing the right not to know (RNTK).

in people making choices that do not reflect their true values and preferences, thus erroneously or accidentally not receiving potentially lifesaving information (Figure 12.1).

Moral Distress

Having examined the full range of effects that honoring or not honoring the RNTK would have on individual patients or research subjects, I turn to an examination of other relevant considerations; namely, those raised by the interests of medical professionals. It is a vexing problem to possess genetic information that one deems to be clinically important, but to be precluded from disclosing it because a patient has exercised their RNTK. These medical professionals are apt to experience what we can colloquially call the “I-can’t-sleep-at-night” problem. More technically, they experience a phenomenon known as moral distress.

Moral distress refers to the situation where one knows the morally correct course of action but is constrained from taking it (Ulrich and Grady 2018). Unlike a classic ethical dilemma, where there are two ethically justifiable, but nonoptimal choices, moral distress involves feeling like there is a clearly correct, but unavailable choice to make. In normal clinical care, moral distress can be found in a range of situations where structural, legal, or institutional barriers prevent someone from doing what they feel would be right.

Given that a patient’s exercise of their RNTK presents a potential risk to medical professionals, the question then is: How much should we weight this concern? Stated another way, when is it permissible for a doctor’s moral interests (i.e., an orientation toward trying to prevent or ameliorate disease) to override patient autonomy? This notion of beneficence trumping autonomy

has been a frequent topic of exploration in the bioethics literature, with some commentators arguing that while autonomy is certainly an important principle, beneficence and autonomy should be complementary. Physician autonomy and morality should also be respected, which sometimes makes it permissible to violate a patient's autonomy.

This is not to say that a medical professional's interests generally, and moral distress in particular, are sufficiently weighty to carry an argument against the RNTK. But considered in the overall context of a rigorous debate about whether or not we should honor an individual's RNTK important medical information about him or herself, it certainly seems like moral distress is at least another relevant consideration in favor of arguing that it is appropriate to be skeptical about a broad, strong RNTK.

Genetic Exceptionalism

It has been popular to argue that genetic information requires special treatment, such as extra privacy protections, enhanced pretest education, and a distinct informed consent process (Green and Botkin 2003). This position was supported by the strongly held notion that there is something different about genetic information. Specifically, genetic exceptionalists have argued that genetic information is often predictive, rather than diagnostic, and thus can be used to predict an individual's future health in ways that other kinds of nongenetic medical information cannot. Genetic exceptionalists have also focused on the fact that since genetic information is an immutable part of your identity and cannot be altered, we should be careful to guard against the psychosocial and economic effects of disclosing genetic risk information. Finally, genetic information has implications for third parties: any genetic diagnosis or risk information is not simply relevant to the patient, but also to their blood relatives.

Nevertheless, as the field of medical genetics has evolved, genetic exceptionalism has been subject to significant criticism. Accompanying this sort of view has been an increasingly powerful chorus of arguments refuting the basic claims of genetic exceptionalists. While genetic information can often predict distant future health (sometimes with high accuracy), there are many examples of nongenetic health information possessing comparable predictive power.

This strong refutation of genetic exceptionalism is relevant to the RNTK debate. Proponents of the RNTK are effectively arguing that the return of any genetic information requires explicitly soliciting patient consent. Since it is standard practice in many clinical situations to disclose certain kinds of nongenomic medical findings without asking for explicit permission, it seems fair to ask whether this instance of genetic exceptionalism is warranted.

Autonomy is obviously an important value in medical ethics; modern social norms have clearly and enthusiastically moved away from medicine's paternalistic history. However, it is not true that patients are asked to make decisions

about every single aspect of their health care. If a patient undergoes a specifically indicated scan, but that scan incidentally reveals a potentially cancerous tumor, a doctor is not going to ask the patient whether they want to learn about the unexpected but important result. Similarly, if a patient receives a routine blood panel to check for a specific indication but the panel returns a panic value indicating a serious acute problem, the physician is not going to ask before disclosing this urgent finding.

These analogies are not perfect. In general, genomic findings are not associated with conditions that require immediate attention, and genetic predisposition is not always equivalent to a diagnosis of manifested disease. The question, however, is not whether genetic information is precisely analogous to the urgent cases presented above. Rather, the relevant question should be whether and why the kind of important genomic information being discussed here warrants special treatment. Given the thorough rejection of genetic exceptionalism, the burden of proof lies with RNTK proponents to make that case.

Conclusion

The currently prevailing view about the RNTK involves an almost exclusive focus on the principle of autonomy. This pure autonomy view results in an environment where individual preferences must be actively sought and respected. At the other end of the spectrum, one can imagine an argument that completely relies on beneficence, justifying forced provision of genetic information whenever it could provide medical benefit to a given individual. In between, there seems to be a more centrist, qualified disclosure view. Embracing libertarian paternalism, we could create a default package of recommended variants to disclose and give patients a choice not to receive genetic information (even if that decision seems objectively unreasonable). This would function as a form of soft paternalism, helping to frame decision making in a way that is thought to lead to more beneficial choices.

I reject the pure autonomy view for the reasons explored throughout this article. I cannot, however, endorse a pure beneficence view either. It seems too paternalistic to force information on someone who is deliberately trying to remain ignorant. Libertarian paternalism is attractive, but partially fails because of concerns about our ability to accurately assess individual preferences for such a complex question. My view falls somewhere between the liberal paternalism and pure beneficence views. For high-impact genetic information, I think that it is a mistake to actively solicit preferences. Instead, patients should be informed that there is a default set of high-impact incidental findings that will be sought and returned. In the rare case that someone independently requests to not learn about this information, in-depth counseling should be provided to ensure that they fully understand the choice being made, but ultimately their decision to remain ignorant should be honored if not knowing consistently remains their clearly stated preference. In short, for

high-impact genetic information, any deviation from regular disclosure should be a clearly defined exception, rather than the basis for a broadly applied conception of the RNTK.

This approach should be relatively uncontroversial for the vast majority of people since most autonomous adults would want to know life-saving information. There are, however, a few predictable exceptions that should be fairly easy to anticipate and accommodate; namely, terminally ill patients, elderly individuals, people with religious objections to treatment, or people in low-resource settings where medical care is not available. These are all cases where clinical action is less certain, so it might be appropriate for medical providers to actively solicit preferences. These kinds of cases represent an important exception to my proposed approach, but I do not believe that we should institute a strong RNTK policy based on a small group that is relatively easy to bracket. The RNTK has become an ingrained part of our lexicon, and though I ultimately believe that we should abandon the term altogether, I recognize that this is unlikely. At the very least, a compelling case can be made that we should at least stop talking about the RNTK in such strong terms.

As a final note, the RNTK debate raises an interesting question about when it is appropriate to utilize institutional power to reduce instances of deliberate ignorance, a topic that is explored in more detail by Teichman et al. (this volume). As mentioned above, it is ironic that a policy to reduce deliberate ignorance of patients by de-emphasizing the RNTK necessarily involves increasing the deliberate ignorance of researchers and clinicians; it might be the case that deliberate ignorance is sometimes a zero-sum game. Furthermore, the framing of a specific deliberate ignorance problem is important. As is made clear by contemporary RNTK advocates, autonomy is often viewed as the controlling principle, with the implication being that people's choices should always be honored. As my preceding analysis hopefully demonstrates, by shifting the frame we can see that there can be cases where we think that people are systematically making less than optimal decisions to remain ignorant. In such cases, where the value of the information clearly outweighs the risks, and where psychological processes are likely to cause some people to make poor choices, it becomes ethically defensible to de-emphasize autonomy in favor of beneficence. In cases like these, it can be justified to use institutional power to create policies or defaults that aim to mitigate the potential for suspect instances of deliberate ignorance.

Acknowledgments

The opinions expressed herein are the author's own and do not reflect the policies and positions of the National Institutes of Health, the U.S. Public Health Service, or the U.S. Department of Health and Human Services. This research was supported by the Intramural Research Program of the National Human Genome Research Institute, National Institutes of Health. This chapter is adapted with permission from Berkman (2014).