Social justice, equality and primary care: (how) can 'big data' help?

Kristin Voigt (kristin.voigt@mcgill.ca)

Final version forthcoming in Philosophy & Technology.

1 Introduction

While there is a growing emphasis among population health researchers on the 'social determinants of health' — i.e. the ways in which various social and economic factors shape health outcomes — it is not clear how primary care systems should respond to our increasing knowledge about how the social determinants of health affect the health outcomes faced by individuals. On the one hand, advocates for primary health care emphasise the importance of primary care for equality in health (see Rasanathan et al., 2011); from this perspective, it seems clear that the primary care system should take into account insights about how various social factors shape the health of individuals (Furler & Palmer, 2010; Institute of Medicine, 2015a). Proponents of patient-centred care and personalised medicine frameworks have also emphasised the importance of addressing the social determinants of health within the primary care system (Garg et al., 2013) and the possibilities offered by big data in this regard (Murdoch & Detsky, 2013). At the same time, it is far from clear how exactly the primary care system should respond to the social determinants of health and, more broadly, to the existence of social inequalities in health among patients, and how a commitment to address the social determinants relates to requirements of medical fairness (Hurst, 2009; Bærøe & Bringedal, 2011). In this paper, I want to shed light on this question by focusing on the role that 'big data' might play in allowing primary care providers to respond to the social determinants that affect their patients' health. Despite growing attention on the role that big data could play across a number of domains, including the health sector,[1] this potential use of data has not been examined in much detail.[2] I focus on the proposal that data from different sources could be integrated in patients' electronic health records, allowing providers to take this information into account when interacting with their patients. This general idea has been proposed and endorsed by the Institute of Medicine (2015a, 2015b) and further refined as part

---

[1] See, for example, Murdoch (2013).

[2] For some discussion of these issues, see Comer et al. (2011) and Stevens (2014).

of the ‘Community Vital Signs’ framework developed by Bazemore et al. (2016) and Hughes et al. (2016).[3]

Of particular concern in this paper is the underlying idea that big data could allow the primary care system to make a contribution to the reduction of social inequalities in health: giving providers access to information about a number of features of individual patients’ environment and circumstances enables them to take into account the ways in which social factors shape their patients’ health outcomes. The idea that primary care could use big data to contribute to the reduction of social inequalities in health is particularly interesting in light of concerns that big data might not be equally beneficial for people from different backgrounds and thus further entrench existing inequalities.[4] One of the aims of this paper is to identify and discuss a number of concerns and tensions arising in connection with the Community Vital Signs proposal — concerns and tensions that also resonate beyond this specific proposal to other interventions that draw on data to improve health outcomes, or indeed outcomes in other areas. At various points the discussion also sheds light on issues that arise more broadly with attempts to integrate this kind of information in a primary care setting even if they do not rely on data in this way — for example, physicians may rely on informal conversations with the patient or on their familiarity with the health centre’s catchment area as sources of information about various social factors.

The paper begins by introducing the Community Vital Signs framework and its central idea of integrating indicators of individual patients’ neighbourhoods into their electronic health record so that primary care providers can take this information into account when interacting with, and providing care to, their patients. I then explore three issues arising from this proposal. First, while questions of privacy have been central to discussions about big data, particularly in the health context, where we are often dealing with sensitive information, the Community Vital Signs proposal also allows us to see that there might be costs to *not* making certain information available. Second, I consider some of the questions arising for primary care from the influence of social factors on health outcomes: how precisely should the primary care system respond to the social factors that contribute to social inequalities in health? Answering this question is crucial in determining how primary care providers should

---

[3] I treat the two papers cited here as developing the same proposal because there is significant overlap between the groups of authors of the two papers: all the authors of Hughes et al. (2016) are also authors on Bazemore et al. (2016); an additional five authors contributed to Bazemore et al. (2016). Bazemore et al. (2016), which is more detailed than Hughes et al. (2016), is the primary source of information for this paper.

[4] For example, because in certain ways of generating data (such as social media), disadvantaged populations are likely to be underrepresented (Hargittai, 2015).

respond to the information that becomes available to them in the Community Vital Signs proposal. Finally, I address problems arising from the use of population level data when dealing with individuals, which creates specific challenges in the context of a proposal such as that proposed by Bazemore et al. and Hughes et al.

2 Primary care and (health) inequality

Over the past few decades, there has been a growing emphasis on the importance of social and economic factors in shaping individuals' health outcomes. The 'social determinants of health' are generally taken to include 'the circumstances in which people are born, grow up, live, work and age, and the systems put in place to deal with illness. These circumstances are in turn shaped by a wider set of forces: economics, social policies, and politics.'[5] Empirical studies have investigated the correlation between various social factors and health outcomes.[6]

Against this background, the role of the health care system has been reexamined.[7] While financial barriers to accessing medical care may seem like an obvious source of social inequalities in health, social gradients in health exist even in countries with universal health care. Indeed, the 1980 Black Report, which generated a lot of the empirical research on the social determinants of health and their impact on health inequalities in the UK, started from the question of why social inequalities in health had persisted despite the introduction of the National Health Service some 25 years prior to the Report (Marmot & Kogevinas, 1987). Subsequent research has confirmed that significant social inequalities in health continue to exist even in universal health care systems: formally equal access to health care, then, is not sufficient to address the social gradient in health outcomes.

In fact, some researchers have expressed concern that some aspects of universal health care may in fact *increase* health inequalities. For example, those with higher incomes and education levels tend to access and receive specialist care to greater extent than those with lower incomes and education levels (e.g. Veugelers & Yip, 2003; see also Starfield et al., 2012: 92). More generally, certain aspects of the health care system may make it easier for better-off patents to benefit than it is for patients from disadvantaged groups.[8] For example,

---

[5] http://www.who.int/social_determinants/thecommission/finalreport/key_concepts/en/

[6] See Braveman et al. (2011) for a review of the available evidence.

[7] While some definitions of the social determinants of health include the health care system as one such determinant (e.g. World Health Organization, 2008: 8), other accounts regard the social determinants as factors that fall outside the health care system (e.g. Sreenivasan, 2008).

[8] See also Starfield et al. (2012) for more general discussion. While my focus in this paper is on industrialised countries, concerns about possible increases in health inequality arising from the

differences in social status between clinicians and patients from disadvantaged backgrounds may lead to worse outcomes for such patients (Thornton et al., 2011), and physicians may be subject to implicit biases that affect their interactions with patients from different groups (FitzGerald & Hurst, 2017).

This raises the question of whether health care systems can, and should, play a role in *reducing* the social gradient in health outcomes, and if so, how. In particular, while we are gaining more knowledge about the importance of social factors in shaping people's health outcomes and their ability to benefit from the health care system, it is far from clear how clinicians should respond to these insights and what role they might have when it comes to the socioeconomic inequalities that affect the patients they see. Should they adapt their treatment to take into account the social factors affecting individual patients? Should they prioritise or spend more resources on patients from lower socioeconomic groups?[9] It is against the background of this disconnect between, on the one hand, our increasing knowledge about the social determinants of health and, on the other, explicit discussion of the role of the primary care system that the proposal discussed here becomes particularly salient: how, if at all, should clinicians respond to the social factors that have such a substantial influence on the health of the patients they see?

3 The Community Vital Signs proposal

The basic idea underlying Bazemore et al. and Hughes et al.'s proposal is simple: when a patient enters primary care, the address they provide can be used to gather information from various data sources about the patient's immediate neighbourhood; these 'community vital signs' can then be included in the patient's electronic health record, along with their medical history, and be made available to the clinician. This process is facilitated by the fact that an address can be linked to different datasets (often datasets that are publicly available) to generate information about the environment immediately surrounding the address.

Bazemore et al. propose a range of indicators that could be selected for inclusion in patients' electronic health records, ranging from neighbourhood economic conditions (e.g. Gini coefficient to capture income inequality), neighbourhood resources (e.g. availability of supermarkets) and social deprivation indices. The underlying rationale for this approach, Bazemore et al. explain, is that

---

introduction of universal health care have also been raised in the context of low-income countries (e.g. Gwatkin & Ergo, 2011).

[9] These issues are discussed further by Hurst (2009) and Bærøe and Bingedal (2011).

> all healthcare professionals could see a patient knowing not just her blood pressure, pulse, respiratory rate, and temperature, but also whether she lives in the presence of poverty, healthy food and water sources, walkable street and parks, and has social capital – or how these add up to predict increased risk of morbidity, early mortality, or other adverse health outcomes. (Bazemore *et al.*, 2016: 408)

Currently, clinicians generally do not have access to this kind of information, unless they obtain it informally through interactions with their patients, or through informal knowledge they might have of the neighbourhood in which the patient lives. This information, Bazemore et al. argue, will allow clinicians to provide better care to their patients: 'Incorporating Community VS [vital signs] into every patient's EHR [electronic health record] will give patients and healthcare providers information that better enables context-informed and community-oriented care' (Bazemore *et al.*, 2016: 410) — 'care that meaningfully accounts for neighbourhood factors that affect patients' health' (Bazemore *et al.*, 2016: 407).

4 Privacy and the costs of (non-)disclosure

The first set of issues raised by this proposal revolves around patient privacy. Privacy is raised as a crucial concern in many contexts where big data is stored and/or used: with considerable quantities of personal information at stake, much of which is highly sensitive and generated in ways that individuals are not always aware of, concerns about privacy have been paramount in the debate.[10] In Bazemore et al. and Hughes et al.'s proposal, address information is linked with information that is not in itself sensitive and often publicly available information. The data that is linked to individual patients and added to their electronic health care records may seem much less sensitive than information about individuals' health; the fact that this information is collected and stored may therefore seem less problematic than it does in other areas. Nonetheless, the proposal may generate its own problems around privacy: the information that the 'community vital signs' proposal would like to see included in individual health care records is potentially very sensitive to patients: poverty and disadvantage is not something that patients are necessarily going to be comfortable revealing about themselves.[11] Poverty can also be a source of shame for individuals. Importantly, patients may also — rightly or wrongly — worry that this

---

[10] See Nuffield Council on Bioethics (2015) for discussion.

[11] The Institute of Medicine also notes that information about social factors may be sensitive to patients. They suggest that providers may therefore want to treat this information differently from other kinds of information but also caution that this approach might perpetuate stigma surrounding particular social factors (Institute of Medicine, 2015b: 332).

information could expose them to bias and disrespectful treatment on part of health care professionals. (I return to this concern in section 5.)

Bazemore et al. do not indicate whether they envisage that patients would be asked whether or not they consent to information being generated from their address and shared with health care professionals. In principle, there is of course a case for individuals having control over information that is potentially sensitive – this means that there are good reasons for letting patients choose whether or not they want this information disclosed to health care professionals. This means that there is a case for building consent procedures into the proposal, so as to make patients aware of the information that is being generated from their address, that it will be included in their record and that it will be available to doctors and nurses who treat them.

However, there might also be costs to *not* having this information available to health care providers; this means that giving patients the opportunity to opt out of the scheme may be costly to them. First, making this information available to providers relieves disadvantaged patients of the burden of sharing this information with providers themselves. There can be costs to patients if they have to go through the process of revealing information about themselves that they are concerned will expose them to the judgement of others, or that they find embarrassing or shameful. The proposal involves indicators of poverty and disadvantage, which are heavily stigmatised. In a somewhat different context, Jonathan Wolff has referred to concerns about having to reveal this kind of information as one of 'shameful revelation': having to reveal to others information that one finds shameful (Wolff, 1998). It might benefit patients to have this information available to health care professionals without patients having to engage in revelations of this kind.

Second, if the alternative is that providers simply make assumptions about their patients' backgrounds, this can also result in distinct costs to the patient. Providers may simply be wrong in the assumptions they make about patients and any adaptations they make to the services they provide or how they communicate with the patient may therefore leave patients worse off than if doctors had access to correct information about their patients' neighbourhood or other aspects of the environment in which they live.

Finally, if we leave it to patients to choose what information to reveal when interacting with providers, this can be problematic because patients may not be best placed to assess the level of disadvantage they face. This concern is not specific to disadvantaged patients; rather, it is a broader concern about how people understand the level of (dis)advantage they face. Such assessments are typically made on the basis of relevant comparisons with other neighbourhoods or environments with which we are familiar. For example, our assessments of the availability of fresh foods in our neighbourhood will be based on how the neighbourhood compares to others we are familiar with.[12] Thus, giving people the option to opt out of a proposal such as Bazemore et al.'s could come at a cost.

While privacy is often considered paramount in debates about data-driven interventions, thinking through the implications of such schemes highlights a number of tensions and complexities around privacy. In coming to an appropriate assessment of these potentially competing considerations will also require a better understanding of various relevant empirical questions, such as whether or not the inclusion of various kinds of information in patients' health care records will increase or decrease any provider bias.

5 Medical fairness and the social determinants of health

Part of the rationale for Bazemore et al's proposal is that the inclusion of 'community vital signs' signals to providers that it is part of their role to take into account this information, just like information about patients' medical history, their blood pressure, previous test results, etc. Providers can then adapt the care they provide to ensure that it is sensitive to the patient's circumstances.

But how, exactly, should this information influence the care health professionals provide? One way to answer this question is to argue that the health care system should, at the very least, not lead to an increase in health inequality and perhaps even seek to contribute to its reduction. As mentioned above, social inequalities in health mean that health needs are greater among disadvantaged groups but it is not clear that they are in fact using health care to a greater extent than other groups, even when the health care system is universal — consider, for example, the data mentioned earlier, which suggested that patients from lower income groups do not use specialists as much as better-off patients do. These kinds of mechanisms can, of course, exacerbate social inequalities in health.

What does this mean for the primary care context? There is, I think, a tension here. On the one hand, the traditional understanding of medical fairness requires that health care providers do not treat people differently on the basis of their social background (for example

---

[12] An interesting example of these kinds of mechanisms is provided by Carla Shedd's (2015) study of students from different Chicago schools, whose assessment of racial injustice varied depending on how much they moved between neighbourhoods with different levels of affluence and diversity: for Black students who did not leave their neighbourhoods to attend school in a more affluent area, frequent negative encounters with police were 'normal' whereas those students who attended schools in more affluent neighbourhoods became aware of differences in police presence and behaviour between different areas.

by prioritising patients from better-off backgrounds) (Hurst, 2009). On the other hand, we also reasonably expect health care professionals to be aware of the context in which patients live, because that affects the conditions with which they present and their ability to adopt the advice given by doctors (e.g. to take time off, exercise more, etc.).

One way to think about this is in terms of outcomes that providers should be striving towards, with support from the community vital signs. One (fairly uncontroversial) goal could be to tailor the advice and information that providers give to individual patients, given what providers (due to inclusion of 'community vital signs' in patient records) know about the environment in which the patient lives. For example, recommendations for increased physical activity to patients who live in areas with high crime rates and who cannot afford a gym are not going to be as helpful as the same recommendation for a patient who lives in a very safe area and who can afford a gym membership as well.[13] Providers could use information about the patients' living environment to make their advice more meaningful for patients, taking into account the environment in which they live and what that tells them about the particular challenges these patients might face.

This approach is clearly consistent with considerations of fairness: if providers fail to tailor their advice and recommendations in this way, then it is likely to favour better-off patients, who are often more like the health care professionals themselves, e.g. in terms of the kind of neighbourhood resources they have access to, incomes, etc. Information about the patient's neighbourhood can help providers correct for the ways in which providing identical treatment for everyone is likely to advantage better-off patients over less well-off ones (Bærøe & Bringedal, 2011). This, it seems to me, is in line with the underlying motivation of the Community Vital Signs model: Hughes et al. (2016), for example, envision that the Community Vital Signs would 'influence point-of-care decisions' (p. 419) and 'tailor clinical and community interventions' (p. 420): 'Knowing social risks', they suggest, 'may inform the clinical recommendations physicians offer' (p. 420).

A second, more ambitious – and more controversial – way to approach the information about 'community vital signs' would be to give greater priority and/or more health care resources to patients who are exposed to negative risk factors. For example, providers can provide longer appointments to patients who live in poor neighbourhoods, or schedule more frequent consultations and follow-ups.

---

[13] In fact, ensuring usefulness of the recommendations providers give to their patients is part of what the Institute of Medicine envisages (Institute of Medicine, 2015b).

This may begin to sound like it is in tension with ‘traditional’ understanding of medical fairness and neutrality; again, however, we can think of this as reflecting the understanding that people from disadvantaged background need additional resources or time to benefit as much as better-off patients from primary care. For example, disadvantaged patients often have lower health literacy, which makes it harder for them to describe their symptoms as efficiently as patients with greater health literacy; social distance between patient and provider can make communication less efficient, and building trust across social divides may be time-consuming (Fiscella & Epstein, 2008). This, I think, helps make the case for a more ambitious goal with respect to what providers should do when it comes to their patients’ social backgrounds.

However, these theoretical considerations must not detract from the possibility that making information about patients’ social background explicit to providers could lead to *less* fairness in health care encounters. The literature has raised concerns about various biases on part of health care professionals (e.g. Burgess et al., 2004; FitzGerald & Hurst, 2017) and unequal treatment within the health care system (e.g. Goddard & Smith, 2001; Starfield et al., 2012). Willems et al.’s literature review finds striking differences in how providers communicated and interacted with patients from different class backgrounds:

> Doctors behave differently during consultations with patients from lower social classes. They are less informative with less educated and lower income patients, possibly because they inaccurately assume that these patients are not particularly interested in learning about their health or do not understand this information. Also they tend to vary the amount of showing affect, involving patients in treatment decisions and express concern by their perception of the patients’ interest in or ability to participate in care. Even though physicians often ground their behaviour on these perceptions, this is not always in accordance with reality. (Willems et al., 2005)

Attempts to include information about the social determinants of health into patients’ medical records should therefore be accompanied by an assessment of whether or not this makes patients from disadvantaged backgrounds more susceptible to mechanisms of this kind.

Thus, if big data facilitate the inclusion of information about social determinants of health into primary care, this raises broader questions about how we think about fairness in the primary care context, how far primary care can and should go when it comes to reducing social inequalities in health and – moving from theoretical to practical questions – whether providing information on patients’ social backgrounds is going to have any impact on possible biases on the part of providers. In this section I suggested that there is a case for a more ambitious interpretation of what health care providers’ role should be when it comes to taking into account the social determinants of health when interacting with their patients but also cautioned that we must assess possible negative effects on disadvantaged patients that might result from making this information available to clinicians.

6 Moving from population level data to individual patients

Primary care providers, such as clinics or practices, may use aggregate data about their patients to gain a better understanding of the population they serve. For example, Hughes et al. suggest that the Community Vital Signs could also be used to improve providers' understanding of the local population and allow them to choose more appropriate priorities (Hughes et al., 2016). This is particularly important because patients are not necessarily drawn from the immediate neighbourhood, and providers' assumptions about the populations they serve may be inaccurate (Bazemore et al., 2011).

However, a central aspect of the Community Vital Signs model proposed by Bazemore et al. and Hughes et al. is that population level data be applied *to individuals:* providers are to use the information they can gather from patients' addresses in their interactions with those patients. Making this move is far from straightforward.[14] Knowing that someone lives in an area that is poor, for example, does not give a provider *certain* information about their individual poverty status or exposure to risk factors. And, equally, someone who lives in a wealthy neighbourhood is not necessarily well-off. The relevant likelihoods (of being poor, for example) may change but that does not allow for definite conclusions about individuals. This concern is particularly well illustrated by a Scottish study on deprivation measures, which found that the majority of low-income people did not in fact live in deprived areas (McLoone, 2001). Area deprivation is often taken as indicator of individual deprivation, which has problematic implications (Katikireddi & Valles, 2015). For providers, then, it might be quite onerous to know what to do with this information. Consider for example the body mass index (BMI). The BMI was explicitly not meant to be used as a diagnostic tool; nonetheless, is commonly used in that way by doctors and health care professionals (Voigt, Nicholls & Williams, 2014). Importantly, this is not a concern about the quality of the data: rather, the very nature of group-based data means that we cannot make certain inferences about individuals.

This problem, of course, does not arise equally with all possible uses of community vital signs. For example, while measures of neighbourhood deprivation do not easily translate

---

[14] See also Sniderman (2015) for discussion of this concern.

into indicators of individual deprivation, indicators of particular neighbourhood characteristics, such as crime rates, can provide important information about the context in which people live. We need to think carefully about how easily this move can be made, and the extent to which different indicators are susceptible to this type of concern.

7 Conclusion

Despite excitement about big data and its potential for improving medical care, there has been relatively little discussion of how data-driven approaches might contribute to the reduction of social inequalities in health. Similarly, while recent debates about primary care have emphasised its potential to address the social determinants of health, much less has been said about how individual providers might respond to the social factors that affect their patients' health. The Community Vital Signs proposal developed by Bazemore et al. and Hughes et al. advances these debates by setting out a concrete and detailed mechanism for integrating data about the social determinants of health into the primary care setting.

In this paper, I focused on three issues. First, given the sensitivity of the information involved, including 'community vital signs' in patients' health care records raises issues of privacy; at the same time, there might be costs to *not* making this information available to providers. Second, I suggested that we need a clearer sense of what goals, precisely, this kind of proposal should pursue: What do we want health care providers to do with information about the communities in which their patients live? While the inclusion of such information in the clinician-patient encounter may challenge traditional conceptions of medical fairness, such inclusion could also enhance fairness, more broadly construed, and help reduce social inequalities in health. Finally, indicators generated at the population level do not easily lend themselves to inferences about individual patients, requiring a more explicit consideration of what the data can, and cannot, tell clinicians about the patients with whom they are interacting. How we respond to these issues and strike a balance between potentially competing considerations will be crucial for determining what kinds of information about their patients it would be useful, and appropriate, for health care professionals to have.

References

Bazemore, A., Phillips, R.L. & Miyoshi, T. 2011. Harnessing geographic information systems (GIS) to enable community-oriented primary care. *Journal of Map & Geography Libraries*. 7(1):71–86.

Bazemore, A.W., Cottrell, E.K., Gold, R., Hughes, L.S., Phillips, R.L., Angier, H., Burdick, T.E., Carrozza, M.A., et al. 2016. "Community vital signs": incorporating geocoded social determinants into electronic records to promote patient and population health. *Journal of the American Medical Informatics Association: JAMIA*. 23(2):407–412.

Braveman, P., Egerter, S. & Williams, D.R. 2011. The social determinants of health: Coming of age. *Annual Review of Public Health*. 32: 381-398.

Burgess, D.J., Fu, S.S. & Van Ryn, M. 2004. Why Do Providers Contribute to Disparities and What Can Be Done About It? *Journal of General Internal Medicine*. 19(11):1154–1159.

Bærøe, K. & Bringedal, B. 2011. Just health: On the conditions for acceptable and unacceptable priority settings with respect to patients' socioeconomic status. *Journal of Medical Ethics*. 37(9):526–529.

Comer, K.F., Grannis, S. & Dixon, B.E. 2011. Incorporating geospatial capacity within clinical data systems to address social determinants of health. *Public Health Reports*. 126:54–61.

Fiscella, K. & Epstein, R. 2008. So Much to Do, So Little Time: Care for the Socially Disadvantaged and the 15-Minute Visit. *Arch Intern Med*. 168(17):1843–1852.

FitzGerald, C., & Hurst, S. 2017. Implicit bias in healthcare professionals: a systematic review. *BMC Medical Ethics*, *18*(19), 71.

Furler, J.S. & Palmer, V.J. 2010. The ethics of everyday practice in primary medical care: responding to social health inequities. *Philosophy, ethics, and humanities in medicine*. 5(1):6.

Garg, A., Jack, B. & Zuckerman, B. 2013. Addressing the Social Determinants of Health Within the Patient-Centered Medical Home: Lessons From Pediatrics. *JAMA: The Journal of the American Medical Association*. 309(19):2001–2002.

Goddard, M. & Smith, P. 2001. Equity of access to health care services: theory and evidence from the UK. *Social Science and Medicine*. 53(9):1149–1162.

Gwatkin, D. R., & Ergo, A. 2011. Universal health coverage: friend or foe of health equity? *The Lancet*. *377*(9784), 2160–2161.

Hargittai, E. 2015. Is Bigger Always Better? Potential Biases of Big Data Derived from Social Network Sites. *The ANNALS of the American Academy of Political and Social Science*. 659(1):63–76.

Hughes, L.S., Phillips, R.L., DeVoe, J.E. & Bazemore, A.W. 2016. Community Vital Signs: Taking the Pulse of the Community While Caring for Patients. *The Journal of the American Board of Family Medicine*. 29(3):419–422.

Hurst, S.A. 2009. Just care: should doctors give priority to patients of low socioeconomic status? *Journal of Medical Ethics*. 35(1):7–11.

Institute of Medicine. 2015a. *Capturing Social and Behavioral Domains in Electronic Health Records: Phase 1*. Washington, D.C.: National Academies Press.

Institute of Medicine. 2015b. *Capturing Social and Behavioral Domains and Measures in Electronic Health Records: Phase 2*. Washington, D.C.: National Academies Press.

Katikireddi, S.V. & Valles, S.A. 2015. Coupled ethical–epistemic analysis of public health research and practice: Categorizing variables to improve population health and equity. *American Journal of Public Health.* 105(1):e36–e42.

Marmot, M.G. & Kogevinas, M. 1987. Social/economic status and disease. *Annual Review of Public Health*. 8:111–135.

McLoone, P. 2001. Targeting deprived areas within small areas in Scotland: population study. *British Medical Journal*. 323(7309):374–375.

Murdoch, T.B. & Detsky, A.S. 2013. The inevitable application of big data to health care. *Journal of the American Medical Association*. 309(13):1351–1352.

Nuffield Council on Bioethics. 2015. *The collection, linking and use of data in biomedical research and health care: ethical issues*. London: Nuffield Council on Bioethics.

Rasanathan, K., Montesinos, E.V., Matheson, D., Etienne, C. & Evans, T. 2011. Primary health care and the social determinants of health: essential and complementary approaches for reducing inequities in health. *Journal of Epidemiology and Community Health*. 65(8):656–660.

Shedd, C. 2015. *Unequal City*. New York: Russell Sage Foundation.

Sniderman, A. 2015. The Role of Physicians in the Era of Predictive Analytics. *Journal of the American Medical Association*. 314(1):25–26.

Sreenivasan, G. 2008. Justice, Inequality, and Health. In *The Stanford Encyclopedia of Philosophy*. E. Zalta, Ed.

Starfield, B., Gérvas, J., & Mangin, D. (2012). Clinical Care and Health Disparities. *Annual Review of Public Health*, *33*(1), 89–106.

Stevens, G.D. 2014. Addressing Social Determinants of Health Using Big Data. In *Big Data and Health Analytics*. K. Marconi & H. Lehmann, Eds. Boca Raton: Taylor & Francis. 105–126.

Thornton, R. L. J., Powe, N. R., Roter, D., & Cooper, L. A. 2011. Patient-physician social concordance, medical visit communication and patients' perceptions of health care quality. *Patient Education and Counseling*, *85*(3), e201–e208.

Veugelers, P. & Yip, A. 2003. Socioeconomic disparities in health care use: Does universal coverage reduce inequalities in health? *Journal of Epidemiology and Community Health*. 57(6):424–428.

Voigt, K., Nicholls, S.G. & Williams, G. 2014. *Childhood Obesity: Ethical and Policy Issues*. Oxford University Press.

Willems, S., De Maesschalck, S., Deveugele, M., Derese, A. & De Maeseneer, J. 2005. Socio-economic status of the patient and doctor–patient communication: does it make a difference? *Patient Education and Counseling*. 56(2):139–146.

Wolff, J. 1998. Fairness, Respect, and the Egalitarian Ethos. *Philosophy and Public Affairs*. 27(2):97–122.

World Health Organization. 2008. *Closing the gap in a generation: health equity through action on the social determinants of health*. Geneva: World Health Organization.