

Moral Cultivation and the Quantified Self: Assessing the Self Understanding of Data Profiles Generated by AI with a Virtue Ethics Approach

La culture morale et le soi quantifié : évaluer la compréhension de soi des profils de données générés par l'IA à l'aide d'une approche d'éthique de la vertu

Ephraim Barrera

Volume 3, numéro 1, 2022

Normativité et intelligence artificielle
Artificial Intelligence and Normativity

URI : <https://id.erudit.org/iderudit/1098932ar>
DOI : <https://doi.org/10.7202/1098932ar>

[Aller au sommaire du numéro](#)

Éditeur(s)

UQAM Département des sciences juridiques
UQAM Faculté de science politique et de droit

ISSN

2563-9250 (numérique)

[Découvrir la revue](#)

Citer cet article

Barrera, E. (2022). Moral Cultivation and the Quantified Self: Assessing the Self Understanding of Data Profiles Generated by AI with a Virtue Ethics Approach. *Communitas*, 3(1), 76–102. <https://doi.org/10.7202/1098932ar>

Résumé de l'article

Les approches qui promeuvent la collecte et l'analyse de données personnelles estiment que les profils de données (data profiles) générés à partir d'algorithme d'intelligence artificielle (IA) représentent une fin souhaitable au profit du « soi quantifié » (ou « automesure connectée », quantified self en anglais). Les tenants du soi quantifié postulent que les profils de données générés par l'IA représentent les vies individuelles de manière objective et authentique. Ils postulent également que le soi quantifié permet une prospérité humaine en offrant aux individus des informations sur leur vie, qui sont basées sur des données concrètes. Tout d'abord, je vais retracer les origines technologiques du soi quantifié. Deuxièmement, le premier postulat sera critiqué en démontrant que le soi quantifié présente une image réduite et subjectivement abstraite de la vie humaine. Troisièmement, je remettrai en question le deuxième postulat, à partir d'une approche de l'éthique de la vertu, pour démontrer dans quelle mesure le concept réduit d'automesure, lié au soi quantifié, est éloigné du concept de cultivation de l'esprit. Pour finir, j'appliquerai un cadre théorique néo-aristotélicien d'éthique de la vertu pour démontrer que la connaissance de soi visée par le soi quantifié entrave le raisonnement pratique des agents individuels.

© Ephraim Barrera, 2022



Ce document est protégé par la loi sur le droit d'auteur. L'utilisation des services d'Érudit (y compris la reproduction) est assujettie à sa politique d'utilisation que vous pouvez consulter en ligne.

<https://apropos.erudit.org/fr/usagers/politique-dutilisation/>

Moral Cultivation and the Quantified Self: Assessing the Self-Understanding of Data Profiles Generated by AI with a Virtue Ethics Approach

Ephraim Barrera¹

ABSTRACT

Supporters of personal data collection and analysis contend that data profiles generated from AI algorithms represent a desirable pursuit for the *quantified self*. Proponents of the quantified self claim that AI-generated data profiles represent a more objective and truthful account of individual lives. They also argue that the quantified self fosters human flourishing by supplying individuals with data-informed accounts about their lives. First, I will trace the technological origins of the quantified self. Second, the first claim will be critiqued by demonstrating that the quantified self presents a reduced and subjectively abstracted picture of human life. Third, the second claim will be questioned, from a virtue ethics approach, to show how the quantified self's reduced concept of self-examination is detached from self-cultivation. Fourth, a neo-Aristotelian virtue ethics framework will be applied to argue that the self-knowledge sought by the quantified self hinders agents' practical reasoning.

KEYWORDS: AI algorithms, data profiles, virtue ethics, AI and human flourishing

RÉSUMÉ

Les approches qui promeuvent la collecte et l'analyse de données personnelles estiment que les profils de données (*data profiles*) générés à partir d'algorithme d'intelligence artificielle (IA) représentent une fin souhaitable au profit du « soi quantifié » (ou « automesure connectée », *quantified self* en anglais). Les tenants du soi quantifié postulent que les profils de données générés par l'IA représentent les vies individuelles de manière objective et authentique. Ils postulent également que le soi quantifié permet une prospérité humaine en offrant aux individus des informations sur leur vie, qui sont basées sur des données concrètes. Tout d'abord, je vais retracer les origines technologiques du soi quantifié. Deuxièmement, le premier postulat sera critiqué en démontrant que le soi quantifié présente une image réduite et subjectivement abstraite de la vie humaine. Troisièmement, je remettrai en question le deuxième postulat, à partir d'une approche de l'éthique de la vertu, pour démontrer dans quelle mesure le concept réduit d'automesure, lié au soi quantifié, est éloigné du concept de

cultivation de l'esprit. Pour finir, j'appliquerai un cadre théorique néo-aristotélicien d'éthique de la vertu pour démontrer que la connaissance de soi visée par le soi quantifié entrave le raisonnement pratique des agents individuels.

MOTS-CLÉS : algorithmes d'IA, profils de données, éthique de la vertu, IA et prospérité humaine

¹University of Ottawa, Center for Law, Technology and Society

Introduction

The desire to uncover personal meaning in data profiles has fuelled the pursuit for the *quantified self* (Quantified Self Institute [QSI], 2016). As a concept, the quantified self describes an understanding of the self as a data-based profile of which inferences and predictions can and *should* be made. The quest for the quantified self has also birthed an international community of users and makers of self-tracking technologies who share an interest in “self-knowledge through numbers” (Wolf & De Groot, 2020), and who support the right and ability to self-track and learn from one’s own data (Quantified Self [QS], 2022). The movement claims thousands of followers worldwide and hosts hundreds of meetup groups, including an annual conference at which enthusiasts share new methods and means of self-tracking (Vallor, 2016, p. 199). This community has also become a focus for academic research on the cultural effects of new technologies, and how data collection and analysis can be used to produce knowledge for citizens to self-use in their daily lives (Wolf & De Groot, 2020).

In *Technology and the Virtues* (2016), Shannon Vallor explains that the quantified self movement believes that data profiles allow us to “transcend the limitations of philosophical and other qualitative methods of self-examination” (p. 198). In harmony with the assertion that life “doesn’t count unless we can measure it” (Vallor, 2016, p. 198), I take the supporters of the movement to believe that algorithmically generated data profiles represent a more objective and truthful account of individual human lives than philosophical reflection. Following this claim, advocates of the quantified self also argue that the pursuit of autonomic self-knowledge through intensive and long-term data collection provides important insights for self-improvement plans (Moore, 2017, p. 2). Hence, I understand the movement to also believe that the quantified self fosters human flourishing through “new technology-driven habits of self-examination” (Vallor, 2016, p. 196) using data-informed profiles about individual lives. Promoters of the quantified self make both claims to justify the increased collection, storage, and analysis of personal data as part of a “right of access to data” (QS, 2022).

Using a virtue ethics perspective, this paper will critique the notion of the quantified self in the context of the broader movement founded by Gary Wolf and Kevin Kelly in 2007 and continuing to the present day (QSI, 2016; QS 2022). The descriptions of the quantified self’s philosophic assumptions are largely drawn from Vallor’s and Phoebe V. Moore’s understanding of the goals and beliefs of the quantified self movement. The appraisal is divided into four parts. First, to provide context, I will trace the origins of the idea of the quantified self by explaining how AI algorithms are employed to generate data profiles. This will involve examining the increased surveillance by data-collecting technologies, the concept of personal data, developments in Big Data, and the practice of

social media analytics to produce data profiles. Second, the first claim made by proponents of the quantified self will be critiqued using Mireille Hildebrandt's work on big data and other authors to argue that, instead of a more objective and truthful perspective, the quantified self presents a reduced and subjectively abstracted picture of human life. Third, the second claim made by followers of the quantified self will be questioned by drawing from Shannon Vallor's virtue ethics approach to show how the quantified self fosters a reduced notion of self-examination that is distinct from the moral project of self-cultivation. Fourth, Alasdair MacIntyre's neo-Aristotelian virtue ethics framework will be applied to the quantified self to argue that the self-knowledge sought by the quantified self inhibits agents' pursuit of flourishing by hindering their practical reasoning.

1. The Origins of the Quantified Self

1.1. Surveillance Technologies and Personal Data

To understand the sociological context of the quantified self movement, it is worth tracing the technological developments that enabled and fuelled the pursuit of quantified human lives. Novel methods of surveillance using data-collection technologies have led to the emergence of AI-generated data profiles which, in turn, motivate the pursuit of the quantified self. Vallor (2016) explains how various means and forms of data surveillance have developed as the digital revolution has taken off (p. 188-189). The significant development in recent years is "The way in which massively networked data storage banks and powerful algorithms now allow us to integrate, aggregate, compare, and extrapolate from the output of these diverse and globally distributed surveillance tools" (Vallor, 2016, p. 189). Most of this analyzed data is not produced by the surveillance mechanisms themselves. Instead, diverse forms of data, including personal data on social media, can be used for surveillance with algorithmic AI technologies (Vallor, 2016, p. 189). "Personal data" or "personal information" can be defined as information that specifically links to particular individuals (L. Francis & J. Francis, 2017). Pieces of personal data that our lives generate (e.g., credit card purchases, emails, web searches, photos, social media "likes," health risks, favourite vacation spots) are likely stored somewhere and are sought as material for surveillance simply because of their potential commercial and security value (Vallor, 2016, p. 189). The increasing phenomenon of personal data accumulation and analysis by surveillance technologies to generate data profiles of individual lives has become known as *dataveillance* (Vallor, 2016, p. 189).

Dataveillance is apt to raise privacy concerns when personal data is gathered from individuals without their knowledge or consent. However, this phenomenon is complexified by the fact that individuals sometimes deliberately record and upload personal data about their lives—or the lives of other citizens—

with the *intent* of feeding that data into surveillance technologies. The term *sousveillance* has been used to describe the culture of expanding and reflexive forms of watching and being watched (Vallor, 2016, p. 190). *Sousveillance* includes practices of *personal or self-surveillance* by wearing data-recording devices such as Fitbit or Apple Watch. Data provided on social media platforms is another example of self-surveillance, since users willingly upload personal data to be viewed on their online profile. Vallor (2016) warns that although isolated pieces of personal data might seem trivial when collected, they can be “profoundly revealing of our selves” (p. 189) when aggregated by algorithms. Advocates of the quantified self support dataveillance practices and encourage the culture of *sousveillance* for the sake of feeding larger amounts of varied information into AI algorithms.

Personal data can be collected by AI and used for surveillance almost anywhere through smartphones and social media (Coeckelbergh, 2020, p. 97-98). In *AI Ethics* (2020), Mark Coeckelbergh reveals that people are often unaware that their data is being gathered, or that personal data used in one context could be later used by third parties in another context (p. 98). Despite social media companies issuing privacy statements and asking for users' consent, it is unclear what happens to users' data or even which data is collected (Coeckelbergh, 2020, p. 99). One common practice by social media companies is *data repurposing*: when “data given in one context are then moved to another domain and used for a different purpose” (Coeckelbergh, 2020, p. 99). For example, data repurposing occurs when social media companies sell the personal data of their users to other companies or move that data between different parts of the same company without their users' knowledge and consent. Through data technologies and widespread data repurposing, personal data is collected, transferred, and analyzed at a massive scale.

Although approaches to protecting information privacy often focus on safeguards for personal data, there is also a privacy concern about *non-personal data* or *metadata* (data about personal data) that can be used to make harmful inferences about individuals (L. Francis & J. Francis, 2017). An example of non-personal data that can be employed to generate such inferences is social media metadata. Facebook discovered that the metadata about a user's profile (such as the *amount* of information shared about a political view rather than the view itself) was more useful and predictive in making inferences about that user than their raw personal data (Zuboff, 2019, p. 163). Consequently, the more data uploaded to social media sites, the more surveillance methods can invade users' privacy by collecting both personal data and metadata to make algorithmic inferences about them and others. The accumulation of metadata has accelerated the phenomenon of dataveillance and expanded the information available for the data profiles of the quantified self (Zuboff, 2019, p. 165).

1.2. Developments in Big Data

The production of data profiles associated with the quantified self is also made possible by developments in big data technologies. Broadly, big data allows the large-scale collection and storage of personal and non-personal data which can be analyzed by AI algorithms. Four defining factors of big data have been identified in the literature: volume (large data storage space), velocity (speed of data creation and analysis), variety (various forms of data), and veracity (uncertainty regarding data quality) (Stieglitz et al., 2018, p. 158). In *Big Data* (2014), Mayer-Schönberger and Cukier explain that big data, along with AI, allows new insights to be extracted by applying algorithms to small pieces of information to infer probabilities and create predictions (p. 12). Big data encourages the phenomenon of *datafication*, in which information about everything possible, even things not previously thought of as data, are translated into data for quantification (Mayer-Schönberger & Cukier, 2014, p. 15). Datafication seeks to quantify aspects of an individual previously thought of as unquantifiable. In social media, datafication operates by converting the personal information and metadata of users into quantifiable data points for algorithmic inferences. This practice of *transforming* the previously unquantified into quantified form is important for understanding the claims of the quantified self. Datafication also enables the practice of *predictive analysis*, in which the implicit latent value of data is unlocked through AI algorithms (Mayer-Schönberger & Cukier, 2014, p. 15). Big data allows massive amounts of personal data to be stored and analyzed with AI algorithms to generate predictions about individual lives from inferred probabilities.

Hildebrandt's work on big data provides further insight into the thinking behind the alleged predictive power of big data and datafication. Hildebrandt (2013) explains that big data can be formulated as ' $n = \text{all}$,' in which ' n ' represents a sample size and "all" refers to the entire population (p. 2). ' $N = \text{all}$ ' means that big data supposedly no longer refers to a limited sample size (as was the case in traditional quantitative methods) but rather, refers to *all* the instances of whatever is being investigated (Hildebrandt, 2013, p. 2). In ' $n = \text{all}$,' the sample equals the population and the uncertainty generated from a mere sample size is absent, as the value of ' n ' exponentially increases (Hildebrandt, 2013, p. 6). It is assumed that any lack of precision in predictive algorithms will eventually be corrected by subsequent collections of data—anything hidden from big data analysis is believed to be revealed with more data (Hildebrandt, 2013, p. 6). Hildebrandt explains that developments in big data also sought out to find hidden *meaning* in large sets of information (Hildebrandt, 2013, p. 7). The meaning was inferred through "knowledge discovery in databases" (KDD)—the process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data (Hildebrandt, 2013, p. 4). The support for increased knowledge via big data

inferences invites *further* datafication of phenomena under investigation to prime limitless opportunities to mine data for new patterns (Hildebrandt, 2013, p. 6). In other words, the assumption that inferences about a large sample size can be extended to the entire population of which that sample is a part, along with the belief that increased data collection will continue to generate novel and meaningful inferences, motivates the increased datafication of human life. Big data provides the technology to store and analyze vast quantities of personal information from various sources. This technology both supports and extends the generation of algorithmic inference about individuals that fuels the pursuit of the quantified self.

1.3. Machine Learning and Data Analytics

The quantified self employs AI algorithms which use *machine learning* to analyze data and produce data profiles. Hildebrandt (2013) describes machine learning as the ability for machines to learn with respect to task “T,” performance metric “P,” and type of experience “E” (p. 4). A machine system learns when it can reliably perform at its metric P at task T following experience E (Hildebrandt, 2013, p. 4). With the example of Google’s targeted marketing, AI uses breadcrumbs of data that users leave online to monitor their online behaviour and match advertising campaigns with predictably receptive audiences. Google collects our information not only from one device but from all connected devices with data. Their AI algorithms “use machine learning to watch what we search, when we search it, and the data we input surrounding those pursuits” (Azati, 2021, para. 4). With the collected data, Google uses predictive analytics to predict behaviours based on previous search and buying history. These predictions are then used to identify personalities, match them to ad types, and categorize these personalities into audience clusters called “lookalikes” that comprise people with similar traits or habits (Azati, 2021, para. 5). Google’s categorization process for creating “lookalikes” mirrors how AI algorithms generate the data profiles of individuals from personal social media data.

Coeckelbergh (2020) further explains machine learning in AI as a statistical process with the underlying task of recognizing patterns or rules in data and using them to explain data or make predictions for future data (p. 83). With AI, this process can be performed autonomously—without direct instruction or rules given by the programmer to achieve task “T.” With only the task or objective being provided, machine learning algorithms can detect rules or patterns not specified by the programmer (Coeckelbergh, 2020, p. 84). According to Coeckelbergh (2020), machine-learning software “can adapt its behaviour to better match the requirements of the task” (p. 84). For instance, machine learning can help distinguish spam from significant emails by analyzing large quantities of messages and learning what counts as spam (Coeckelbergh, 2020, p. 84). The starting point

for machine learning algorithms is data rather than theory; the data is “active” in that it defines what the AI does next (Coeckelbergh, 2020, p. 84). Recently, machine learning based on big data has generated significant interest due to the voluminous amount of available data and growth in cheaper computing power (Coeckelbergh, 2020, p. 87). Some researchers have called this phenomenon the “dataquake,” a situation in which we all produce large quantities of data by means of our increasing digital activities (Coeckelbergh, 2020, p. 87). The autonomous pattern-recognizing power of machine learning coupled with the contemporary “dataquake” provide the conditions that stimulate the quantified self.

1.4. Social Media and Data Profiles

Research on social media analytics has revealed that it has evolved to become an “important driver for acquiring and spreading information in different domains” (Stieglitz et al., 2018, p. 156) such as business, entertainment, science, crisis management, and politics. The term “social media big data” has been used in the literature to describe the enormous growth of social media platforms and the increasing accumulation of social media data (Stieglitz et al., 2018, p. 156). Social media offers data in numerous forms such as texts, pictures, videos, sounds, and geolocations. Social media data can also be analyzed to gain insights into trends and influential actors. For instance, research has shown that Twitter data can be analyzed to show how a person's mood changes with the time of day (Stieglitz et al., 2018, p. 156). The field of social media analytics has identified a four-step framework to conduct its analysis: discovery of latent patterns in data, tracking of data, preparation of data, and analysis of data depending on the purpose (Stieglitz et al., 2018, p. 158). Additionally, social media is unique in that it involves content that is not only created by individuals, but is also continuously modified by all users in a participatory and collaborative fashion (Stieglitz et al., 2018, p. 157). Social media big data thereby encourages the culture of *sousveillance* and enables the algorithmic analysis of personal data to make inferences and predictions about users in the quest for the quantified self.

A well-researched example of the intersection between social media data and AI algorithms to generate data profiles can be found in the workplace. Ajunwa and Schlund (2020) describe how machine learning algorithms have promoted the “quantification of the worker in a manner and to a degree, previously unseen in history” (p. 806). One example they point to is the increasing use of automated hiring platforms (AHPs). AHPs now aid companies through the sourcing, screening, interviewing, and selecting stages of the hiring process (Ajunwa & Schlund, 2020, p. 806). AHPs are directed to find attractive candidates using targeted advertising or matching technology that construct predictive models to “generate a pool of jobseekers with predetermined, sought-out

characteristics" (Ajunwa & Schlund, 2020, p. 806). Companies also use AHPs to screen candidates according to their potential based on models that "predict how likely a job candidate is to work well with others or steal from the company" (Ajunwa & Schlund, 2020, p. 807). Recently, AHP vendors have offered social media background checks that analyze an applicant's social media and online history to predict the likelihood that the applicant will engage in toxic workplace behaviours including bullying, harassment, or drug use (Ajunwa & Schlund, 2020, p. 807). These social media background checks can also reveal information about an applicant that employers are not legally permitted to use for hiring decisions (e.g., ethnicity, sexuality, disability, or pregnancy status) (Ajunwa & Schlund, 2020, p. 810). AI algorithms have thus been used to harness personal social media data to create predictive-behavioural models to determine who can be hired and even who is eligible to see certain job opportunities.

1.5. Data Profiles and the Quantified Self

The workplace example helps to highlight how predictive data profiles of individuals are constructed with AI algorithms and social media big data. Ethical concerns are often ignored or belittled in the quest for data profiles due to the promise that such a quantified self is thought to hold. In his article, "Is Human Judgment Necessary?" (2020), Spaulding addresses the subtle but significant ways that AI alters the epistemological field in which human judgment occurs (p. 391). With the example of AI psychological profiling, profiles are built of individual voters, consumers, or possible security threats from "past observable choices (both our own and those of our supposed doppelgängers) to generate a probabilistic assessment of future actions and likely response to future stimuli" (Spaulding, 2020, p. 391). Data profiles are developed from only those past actions which were open and available to recording. The hope is that predictive analytics generated from these profiles will liberate us from biases, errors, and cultural blindness in our judgments (Spaulding, 2020, p. 391). However, Spaulding (2020) points out that this "liberation" is "to a future heavily determined by rationally calculated abstractions aggregated from observable data of our past choices" (p. 391). Our future is reconstrued as a probabilistic calculation of how we would respond in situations based on algorithmic analysis. The past, at least the past that can be quantified, determines our future. With this technology, the predictions of our data profile convey whom we are now and whom we will become.

As these data technologies have developed, the notion of the quantified self—the human person construed as a data profile of which predictive inferences can and ought to be made—has emerged (Wolf, 2018, p. 1). Founded by Wolf and Kelly in 2007, the quantified self movement aims to explore "what new tools of self-tracking are good for" and "to create an environment where this question

can be explored on a human level" (QSI, 2016). The movement's support for "every person's right and ability to learn from their own data" (QS 2022) translates into an implicit support for the increasing datafication and quantification of human life. Philosophically, Vallor (2016) understands the quantified self movement to embrace technologically driven habits of self-examination to attain human flourishing or the "good life" through the quantified self (p. 196). Vallor (2016) describes the concept of the quantified self as a picture of one's own existence that is produced by constant surveillance using biometric and mobile sensors to produce copious quantities of streaming data about one's psychological and physical states and activities (p. 196). Since the movement believes in "self-knowledge through numbers" (QS, 2016), they also hold that the more data available to collect and analyze about our lives, the better we can understand ourselves and improve our flourishing (Vallor, 2016, p. 196). Consequently, the quantified self movement supports the culture of sousveillance's increased use of self-surveillance technologies that utilize mobile or wearable sensors (e.g., Fitbit, smartphone apps, video recording) (Vallor, 2016, p. 198). For adherents of the quantified self, personal data generated from social media is another source of data that ought to be collected by surveillance technologies and analyzed by AI to enhance our self-knowledge through increased datafication.

1.6. Assessing the Quantified Self

The preceding analysis traced the concept of the quantified self from novel surveillance technologies and sousveillance practices fuelled by big data, AI algorithms, and the datafication of human behaviour. Personal data on social media has been increasingly aggregated for purposes such as targeted marketing and workplace hiring decisions. This data is tracked, collected, prepared, and analyzed to produce data profiles of individuals. These profiles are then used to make inferences about individuals' future behaviours and categorize them according to groupings. For the quantified self movement, the growth of self-tracking tools and access to personal data represents progress toward the desirable goal of the quantified self (QS, 2022). Considering these goals, Vallor (2016) pinpoints two claims of the quantified self movement and its supporters that can be identified for appraisal. First, the quantified self purports to provide a more objective and truthful account of individuals, in terms of their behaviour and character, than other methods of assessment (Vallor, 2016, p. 193). This means that the personal data collected and analyzed into data profiles represents a more factual account of whom an individual is and how they are likely to act than 'traditional methods' such as interviewing, self-assessment, or philosophical reflection (Vallor, 2016, p. 195). Second, since data profiles are believed to convey a more objective picture of individual lives, the quantified self claims to foster human flourishing by enabling individuals to better develop

themselves through the insights offered by these profiles (Vallor, 2016, p. 196). Further quantification and datafication of one's life should thus be encouraged on this view for the benefits of enhanced self-examination and of a correspondingly more truthful society (Vallor, 2016, p. 192). These two claims about the quantified self championed by its supporters will now be assessed.

2. The Quantified Self as Reduced and Subjective

2.1. Practical Modelling Versus Truthful Knowing

Instead of providing a more 'objective' or 'truthful' portrayal of human life, the quantified self interprets human beings in a reduced and subjectively abstracted way. The quantified self is informed by the information collected and analyzed by big data. Hildebrandt (2013) argues that talk about 'true' or 'untrue' does not apply to big data, because the technology engages in data modelling rather than truth-telling (p. 2). She contends that the underlying question raised by big data is whether the modelling works for our purposes, and "what kind of humans we will become when interacting with the models that Big Data generates to figure us out" (Hildebrandt, 2013, p. 2). To argue that the quantified self presents a more accurate and truthful account of human life is thus to make a category error. Data profiles generated by big data make inferences based on models that aim to predict X phenomena under Y conditions. The criterion for a "good algorithm" is its effectiveness to predict X when Y arises, not whether the algorithm accurately conveys the full and nuanced truth about X or what X relates to. The distinction between algorithms that *work* (function well in effectively predicting what they set out to predict) and algorithms that are *truthful* (accurately and fully convey a phenomenon) is ignored, or at least diminished, by the quantified self. In a sense, the category error is made by the concepts of "truthful" and "effective" collapsing onto each other when describing an AI algorithm as "accurate." Thus, there is no justification for the claim that the quantified self portrays a more truthful account of human life. Truthful accuracy is not a criterion that determines how algorithms function or how they are assessed as effective. What data and algorithms convey is essentially reduced — a "good algorithm" merely has the ability to effectively predict an aspect of a particular phenomenon in certain circumstances within the limitations of the data that it possesses.

Vallor also questions the alleged accuracy or truthfulness of the quantified self. Drawing from Evgeny Morozov, she argues that those who believe that increased data collection will lead to a more truthful society are guilty of *information reductionism* (Vallor, 2016, p. 192). One can hold to the existence of objective truth while also recognizing that reality is more intricate than data models can portray (Vallor, 2016, p. 192). The moral truth of a situation will always

be richer and more complex than first appearances may suggest and, for Vallor, moral reality “will always overflow any concrete human representation of it” (2016 p. 192). Our response to moral situations ought to include efforts to enhance our perspective and understanding of the situation in order to “see” all of its morally relevant features. To truthfully understand a situation and form good moral judgments according to the situation’s particularities, agents must expand their moral attention to perceive the nuances and complexities of moral reality rather than seek false clarity in data modelling.

Similarly, in *The Voice of Liberal Learning*, Michael Oakeshott argues that information alone “does not indicate how, on any occasion, it should be used” (2001 p. 50). For Oakeshott, “information” is itemized facts and propositions that provide the conditions of performing well and principles that explain performance, whereas “judgment” is the knowing *how* that must be partnered with information. Judgment is also the tacit or implicit component of information that “is unspecifiable in propositions” (Oakeshott, 2001, p. 49). The worldview of the quantified self reduces the intricacy of moral situations by assuming that such circumstances can be comprehensively and truthfully conveyed via data modelling. However, the quantified self ignores, or at least obscures, the consideration that we, and especially AI algorithms, may fail in accurately perceiving moral reality and in judging our information adequately. For Oakeshott and Vallor, moral judgment is a human skill that requires agents to cultivate better ways of *attending* to their circumstances and imagining what to *do* about them. Information alone is insufficient for truthful moral perception or judgment. Hence, our response to the complexity of reality should not simply be to collect more “bits” of morally salient information, as the proponents of the quantified self contend, but to “cultivate better ways of seeing, questioning, thinking about, and listening to it” (Vallor, 2016, p. 192).

Since the interpretive framework for AI algorithms is limited to data patterns, the technology is unable to cultivate these human skills of “seeing” and judging well. Consider the moral complexities in thinking about environmental protections. If a local government were to morally deliberate on whether to place environmental protections on a section of green space that could be used for commercial development, they would need to consider various economic, social, political, scientific, and aesthetic factors. Although *some* of these factors could be given reasonably accurate data points (e.g., certain price valuations, or projected pollution levels), other factors would be essentially *reduced* in the process of quantification (e.g., the aesthetic beauty of the land, or the sociocultural benefits for both protected outdoor spaces and commercial development). Moreover, the various considerations within these factors would have to, in some way, be *rank-ordered* according to their moral value so that a choice could be made between competing considerations. Should the projected economic benefits be given greater moral weight than aesthetic

considerations? How can the socio-cultural benefits of commercial development be measured against the benefits of reduced pollution? The moral values that we could assign to these factors is an ethically complex issue that AI algorithms would be incapable of definitively answering for us in a way that is comprehensively truthful and that regards all of the salient moral elements of the situation.

Advocates of the quantified self can also be charged with informational reductionism for their assumptions about what “truthfulness” entails. The quantified self movement may argue for an idea of “truth” that is conceptually linked to “openness” and a culture where all behaviours are viewable to the public as in futurist David Brin’s book *The Transparent Society* (Vallor, 2016, p. 190). The desire for data openness underpins the quantified self movement’s call for the “right to data” as part of the right to participate in science under Article 27 of the *Universal Declaration of Human Rights* (Wolf, 2018, p. 1). Open-data enthusiasts like former Google CEO Eric Schmidt and David Friedberg, CEO of a big data analytics firm owned by Monsanto, contend that the pursuit of transparency through data is also a social good that will lead to a more “truthful society” (Vallor, 2016, p. 191-192). Since personal data would be collected and analyzed to reveal more about the behaviour of citizens, supporters of the quantified self maintain that the resulting openness and transparency will morally benefit society overall. In response, Vallor (2016) points to the 2015 Ashley Madison hackers as an example of how truth conceived as “absolute openness about information” leads to tragedy (p. 192). The hackers indiscriminately dumped data from the Ashley Madison website to expose an unethical business model and numerous would-be marital cheaters (Vallor, 2016, p. 192). This recklessness endangered the lives of closeted gay and lesbian citizens living under oppressive regimes, provided foreign intelligence agencies with mountains of leverage for blackmail against public officials, and left a trail of devastation in the collateral damage of suicides and family separations (Vallor, 2016, p. 192). For Vallor (2016), this represents the fruits of a “transparent society” without virtues such as humility, care, and wisdom (p. 192). Truthfulness is a character trait exercised by virtuous agents with the moral wisdom to know when and how to reveal sensitive information. The quantified self’s prioritization of openness and transparency is thus a deficient framework for describing what a genuinely truthful and morally wise society should aspire to.

2.2. The Subjective Element of Quantification

Regarding the supposed objectivity of the quantified self, Hildebrandt questions big data’s assumption that ‘n = all’ by arguing that “quantification always implies a preceding qualification” (2013 p. 8). To quantify the flux of life into machine-readable data points involves a series of human decisions on what counts as data, what counts as the same type of data, and what realities fit which

attributes used in the data model (Hildebrandt, 2013, p. 8). All these decisions involve an element of subjective interpretation in the creation and use of algorithmic models. Consequently, there is no such thing as “raw” or objective data; data is constructed by human beings and employed for particular purposes (Hildebrandt, 2013, p. 8). This is not to imply a form of universal relativism that denies the existence of any facts or truth, but rather to point out that there is an element of *personal participation* in the use and application of data models. ‘N’ is never “all” because the flux of life can be translated into machine-readable data points in a *variety* of ways, and the way chosen significantly impacts the outcome of the model (Hildebrandt, 2013, p. 8). Thus, the quantified self cannot assume to hold a more objective portrait of human life. The inferences and categories of the quantified self are constructed by human-made decisions and interpretations about the relevance and weighting of certain quantified realities.

Coeckelbergh also points out the element of human judgment in data by underlining what AI technologies lack. On their own, AI lacks our understanding of relevance, experience, sensitivity, and wisdom - human traits necessary for intelligently using and applying data models (Coeckelbergh, 2020, p. 90). It is commonly known in statistics that correlations do not necessarily imply causal relations; there are correlations in which the variables only *appear* to be causally related and are actually related due to the presence of another invisible or unnoticed factor (Coeckelbergh, 2020, p. 90). Without human direction and interpretation, the inferences made by AI algorithms will be irrelevant or nonsensical. For example, an AI algorithm might find a correlation between the divorce rate in Maine and the per capita consumption of margarine, but humans are needed to decide which correlations deserve further study as potential causal relations (Coeckelbergh, 2020, p. 91). Additionally, creating data sets that inform data profiles involves choices about how to *abstract* from reality. Human decisions are required to determine if certain abstracted data representations are appropriate for a particular purpose. These data abstractions also result in representations of reality rather than reality itself (Coeckelbergh, 2020, p. 91). This is similar to how a map is not itself a territory, but a representation of a territory—about which humans have made choices concerning how to design the map for a particular purpose (e.g., maps for hiking versus maps for car navigation) (Coeckelbergh, 2020, p. 91). The abstractions generated by AI to create data profiles are models of reality rather than reality itself, and they include human choices about how to design the data sets that algorithms train on (Coeckelbergh, 2020, p. 91). As such, the notion that the quantified self is permeated by superior objectivity in the absence of an element of subjective human judgment is an illusion.

3. Assessing the Self-Examination of the Quantified Self

3.1. Vallor's Virtue Ethics Approach

The second claim made by the quantified self-movement is that the pursuit of autonomic self-knowledge provides the knowledge to set out self-improvement plans (Moore, 2017, p. 2). Supporters of the movement claim that improved self-tracking technologies foster human flourishing by allowing us to “transcend the limitations of philosophical and other qualitative methods of self-examination” (Vallor, 2016, p. 198). However, using Vallor's virtue ethics approach, this claim can be shown to be quite dubious given its reduced conception of what self-examination entails. To better understand what we should do with these new data and self-tracking technologies and what they will do *to us*, she offers an ethical strategy for cultivating the type of moral character that can “aid us in coping, and even flourishing, under such challenging conditions” (Vallor, 2016, p. 10). After Vallor's virtue ethics has been outlined, her critique of the quantified self's contribution to flourishing and self-development will be explored.

Vallor assesses the idea of the quantified self through a unified virtue ethics approach drawing from the shared moral commitments in classical Aristotelian, Buddhist, and Confucian virtue traditions. In her view, these three traditions outline a vision of moral life that leads to the cultivation of a virtuous character (Vallor, 2016, p. 32). Ideally, this life begins with proper moral habituation into social roles and responsibilities that, with practice, become increasingly self-directed, reflective, and intelligent (Vallor, 2016, p. 42). Through this process of *moral self-cultivation*, Vallor (2016) explains that agents acquire the virtues — roughly understood as stable habits of mind and action that denote excellence in character (p. 50-51). For example, to be a virtuous agent is to have a character that has cultivated the virtue of courage, which means that the agent is able and willing to act in a courageous way, out of a motive to do what is good, that also accounts for the particularities of the situation. Moral self-cultivation also trains agents to correctly understand and value the ultimate aims of moral living along with the appropriate means to attain those aims in various contexts (Vallor, 2016, p. 42). Moreover, Vallor (2016) distinguishes between “thick” and “thin” moral concepts to identify the shared “thin commitments” between classical Aristotelian, Confucian, and Buddhist traditions. *Thin concepts* only supply the essential structure or skeleton of an idea, whereas *thick concepts* furnish the idea in greater detail (Vallor, 2016, p. 43). Four shared “thin commitments” that Vallor (2016) articulates between these three traditions are a conception of the “highest good” or “human flourishing” that serves as the aim of ethics, of “moral virtues” as cultivated states of character, of the practical path to moral self-cultivation, and of what human beings are generally like (p. 44).

In Vallor's (2016) unified and "thin" conception of what virtue ethics entails, one of the core habits in the practice of moral self-cultivation is *reflective self-examination* (p. 90). This habit is to become a lifelong practice that "aims to discern how well one's actions, feelings, thoughts, and beliefs conform with the moral self to which one aspires" (Vallor, 2016, p. 90). An agent exercising reflective self-examination should foster the virtues of humility and honesty to remain attentive to the particular weaknesses and faults that their character is generally subject to (Vallor, 2016, p. 90). The practice of this habit should also inspire a sense of personal responsibility in the agent to correct their faults and act in a more reliably virtuous manner. Moreover, the practice should "engender appropriate joy in the ongoing experience" (Vallor, 2016, p. 90) of moral self-cultivation and virtuous living. In this unified virtue ethics framework, self-development for human flourishing is attained through moral self-cultivation, and the habit of reflective self-examination helps us to have the self-knowledge required to respond appropriately to new and unfamiliar circumstances — a trait increasingly needed amid the fluctuating conditions of 21st-century life (Vallor, 2016, p. 23).

3.2. Self-Examination and the Cultivated Self

In Vallor's virtue ethics, the pursuit of self-knowledge via reflective self-examination is not an end in itself, but a means to moral self-cultivation or the *cultivated self* (Vallor, 2016, p. 196). She explains that the cultivated self is characterized by an improved, self-conscious, and lifelong effort to bring one's examined thoughts, feelings, and actions nearer to some ethical ideal (Vallor, 2016, p. 196). Related to the self-cultivation that Plato called the "care of the soul," this involves "philosophical habits of self-awareness that enable a gradual realignment of one's actions, values, emotions, and beliefs with the Good" (Vallor, 2016, p. 196). Unfortunately, leading the examined life is a struggle for most since we are often blind to our own failings, unable to distinguish between genuine and false goods, and distracted from virtuous aims by superficial desires. Vallor's solution to these obstacles to moral self-cultivation is reflective self-examination, but it must be carried out wisely (2016 p. 198). We should examine ourselves with flexibility and moderation, and without the obsessiveness that produces paralyzing anxiety or the endless self-critique of narcissistic overestimation (Vallor, 2016, p. 197). Developing this habit is essential to cultivating moral character in novel technological contexts, but it is also vulnerable to disruption, neglect, corruption, and replacement by counterfeits (Vallor, 2016, p. 197). For Vallor (2016), acquiring the cultivated self requires that agents exercise self-examination "as a moral practice" (p. 197).

The pursuit of the cultivated self in virtue ethics also employs "philosophic technologies." Vallor traces the language of "technology" in this sense to the Greek concept of *techne*, a craft in which a product (in this case the self) is

gradually constructed and shaped toward a particular end (2016, p. 197). These technologies are typically associated with philosophical practices including Socrates's dialectical questioning, modelling after exemplary persons, immersion in moral education, narrative, habits of confession, meditation, and reflective examination of conscience (Vallor, 2016, p. 197). However, Vallor states that for these practices to be authentic, "Techniques of self-care must be consciously and reflectively embraced by the moral agent being cultivated" (2016 p. 197). Agents must intentionally and thoughtfully employ philosophical technologies that aid reflective self-examination with the goal of pursuing the cultivated self. By cultivating the virtues and examining one's life through philosophical technologies, agents also move toward a vision of human flourishing classically known as the "good life" (Vallor, 2016, p. 196-197).

Despite some fluidity in the concept of the "good life" or "human flourishing," Vallor identifies three core notions that have endured in the Aristotelian, Confucian, and Buddhist virtue traditions. First, a good human life presupposes my ability to choose to reflect upon and attend to the trajectory of my moral development (Vallor, 2016, p. 199). The self-cultivation characteristic of a good life requires the capacity to assess the direction of my life through different contexts and activities. Second, I must take steps to actively cultivate or steer the trajectory of my life in the right direction (Vallor, 2016, p. 199). The cultivated self involves the deliberate pursuit of moral growth through thoughtfully chosen actions. Third, I have a duty to be concerned with the aspects and activities of my life that are central to my moral character (Vallor, 2016, p. 199). This means that I must seek to cultivate the virtues (e.g., courage, self-control, honesty, prudence) while avoiding the vices (e.g., cowardness, self-indulgence, dishonesty, foolishness). Now that Vallor's conception of virtue and the cultivated self has been outlined, it can be applied to the belief that the pursuit of the quantified self contributes to human flourishing via augmented self-examination. How does the pursuit of self-knowledge in the quantified self compare with these three notions of human flourishing sought by the cultivated self?

3.3. The Quantified Self Versus the Cultivated Self

At first, it might seem that the quantified self aligns excellently with the first notion of the cultivated self due to its obsession with the trajectory of a human life according to tracked data values. Proponents of the quantified self are excited about "finding ways to track ever more personal variables, with even greater degrees of mathematical precision and reliability" (Vallor, 2016, p. 200). However, as the number of tracked variables expands, Vallor suggests that attending to oneself could become more difficult, as many of these diverse variables might not actually be morally meaningful (2016 p. 200). For example, does tracking the details of my watch history, screen time, sleep cycles, driving habits, web

searches, and walk patterns really serve to help me identify my moral trajectory? Self-cultivation requires moral attention, which is about *screening out* information as much as it is about taking information in (Vallor, 2016, p. 200). It is doubtful that self-cultivation can be served when there is an ever-expanding range of self-surveillance, rather than a thoughtful selection of and attention to the most salient variables for moral growth (Vallor, 2016, p. 200). Tracking excessive or irrelevant variables is akin to the hyperactive and overly self-conscious efforts at self-cultivation (Vallor, 2016, p. 200). The vision of the quantified self also emphasizes the increasing number of tracked variables without providing a clear sense of how these variables are to be rank-ordered in terms of their moral importance for a cultivated life. Thus, the quantified self fails in the first notion of the cultivated self. The plethora of variables and traits tracked by the quantified self fosters an informational overload that hinders proper attention to the salient features in an agent's life that reveal their moral trajectory.

How does the quantified self hold up to the second notion—the deliberate pursuit of moral growth? Vallor points out that few of the enthusiasts of the quantified self “express a clear sense of what kind of self they wish to cultivate overall” (2016 p. 201). The quantified self movement seems mostly concerned with information and knowledge for their own sake, rather than for moral cultivation (Vallor, 2016, p. 201). Although many in the movement seek to track certain variables precisely because they could provide “personal meaning” (QS, 2022) (e.g., snoring data could help provide insight into sleep quality), they fail to give an account of how information about personal states can and *should* be used for moral development. The kind of meaning sought and ends pursued (if any) with the data of the quantified self is highly ambiguous (Vallor, 2016, p. 201). That said, Vallor does allow that the techniques of the quantified self *could* someday be recruited for self-cultivation, although this does not seem to be the movement's goal at this point (2016 p. 203). This vagueness hinders agents' ability to judge wisely by choosing actions that will contribute to their moral cultivation.

In my view, the lack of a purpose or *telos* connected to self-cultivation in the movement is no small matter. To encourage “self-knowledge through numbers” and support the increased datafication of human life is ultimately to advance values and policies that will change society to accord with the vision of data technologists. If private and public sector actors are given the freedom and funding to increasingly collect data and use AI algorithms in generating data profiles and predictive analyses about citizens without moral guidance and limitations, the results could be disastrous for human flourishing. The sobering reality is that these technologies could be used (intentionally or unintentionally) to perpetuate inequalities and discrimination toward the disadvantaged, while benefitting the rich and powerful. Thus, it is morally irresponsible for supporters of the quantified self to encourage increased datafication and algorithmic applications without clear moral guidance to ensure that these practices will be

wisely limited and properly directed to serve the ends of moral self-cultivation and human flourishing.

Regarding the third notion of the good life, do the data profiles of quantified selves clearly identify the virtues and vices of human character and the moral features of an agent's activity? It is doubtful that the moral dimensions of an agent could be reduced to a tidy list of variables with discrete numeric values, akin to how we can quantify variables such as body weight, vitamin consumption, and spending activity. Vallor (2016) maintains that the "moral dimensions of the self are among the most difficult to translate into numbers" (p. 201). Enthusiasts of the quantified self readily assign numbers to philosophically complex concepts such as "happiness" and the "good life," but they do so by defining these concepts in terms of a subjective sense of well-being and engagement (Vallor, 2016, p. 201). As a result, a high value for "happiness" could be assigned to agents with a strong subjective sense of contentment, who nonetheless have a vicious character and poor self-knowledge. At most, it seems that the goals and methods of the quantified self would provide an *alternative* to the philosophical technologies of the self that could be used as a supplementary tool to cultivate virtue and promote flourishing (Vallor, 2016, p. 201). However, Vallor points out that the quantified self movement believes that these philosophical technologies are destined to be *replaced* with the more "objective" and "scientific" data tools, and they support increasing reliance on technological devices to record and analyze personal data for enhanced self-examination (2016, p. 202). As explained in the previous section, the movement erroneously believes that the data profiles generated by these devices will reveal a more truthful account of the agent than the philosophical practice of self-examination. Hence, the quantified self fails in the third notion of the good life by making a category mistake. Even the most accurate and comprehensive recording of your past and present states would not constitute an examined life, because "a dataset is not a life at all" (Vallor, 2016, p. 202).

In the classical virtue ethics traditions of Aristotelianism, Confucianism, and Buddhism, our life includes our future, because the examined life is construed as a project of ongoing self-cultivation, rather than a finished achievement. Unlike the quantified self movement, self-knowledge via reflective self-examination is not prized in these traditions for the data it reveals but "for the transformative nature of the practice itself and the dignity it confers upon those who take it up" (Vallor, 2016, p. 202). The examined life is worth living because it embodies chosen habits of mind and action that constitute a person who takes up responsibility for their own moral being as a moral agent (Vallor, 2016). The quantified self movement does not seem to clearly identify and pursue moral self-cultivation or self-development in their "self-knowledge through numbers" (QS, 2022). The deliberate cultivation of moral virtues seems absent from or irrelevant to the ends sought through self-quantification. Consequently, the algorithmically examined

life provided by the data profiles of the quantified self does not necessarily contribute to human flourishing.

4. Practical Reasoning and the Quantified Self

4.1. MacIntyre's Neo-Aristotelian Human Flourishing

Whereas Vallor's virtue ethics approach is useful for questioning the claim that the quantified self contributes to flourishing through enhanced self-examination, MacIntyre's virtue ethics approach is insightful for further showing how the self-knowledge encouraged by the quantified self *hinders* agents' pursuit of flourishing. Rather than synthesizing the "thin" moral commitments from multiple classical traditions as Vallor does, MacIntyre's work advances and extends a neo-Aristotelian framework for human flourishing. MacIntyre's neo-Aristotelianism draws from various thinkers including Sophocles, Aquinas, and Marx to present an account of virtue ethics for 21st-century life. After the core components of his view have been outlined, particular attention will be paid to his notion of *practical rationality* to further assess the quantified self. Applying MacIntyre's neo-Aristotelian virtue ethics reveals that the pursuit of the quantified self hinders human flourishing by inhibiting agents' exercise of practical rationality.

MacIntyre conceptualizes neo-Aristotelian human flourishing in terms of the teleological scheme proposed by Aristotle and developed by Aquinas. Aristotle understands human flourishing as *eudemonia* (often translated as "blessedness, happiness, or prosperity"), roughly understood as a state of well-being and doing well in being well in one's life activity (MacIntyre, 2007, p. 148). To flourish as a human being is to attain this state of activity and thereby achieve one's *telos*, our purpose or "final end" as human beings. The virtues are the qualities of mind and character that help us to flourish and achieve our final end, whereas the vices are the qualities that frustrate our flourishing (MacIntyre, 2007, p. 148). Aristotle characterizes the virtues as dispositions to act and feel in particular ways, so to act virtuously is to act from an inclination formed by habitual cultivation (MacIntyre, 2007, p. 53). Human beings are understood as rational animals who can train their desires, habits, and rational agency to achieve their final end. Cultivating the virtues also helps agents to develop their *practical reasoning*, an ability needed to identify what goods are at stake in their specific circumstances and how they should act for the sake of the good (MacIntyre, 2016, p. 162). As MacIntyre (2007) explains, Aristotle's precepts for good action are connected to certain prohibited vices and encouraged virtues that enable us to move from potentiality to act, to realize our true essence as rational animals, and to flourish as human beings (p. 53).

In Aquinas's view of human flourishing, the final end is an unqualified good that "stands to those other goods as a measure stands to what is measured" (MacIntyre, 2016, p. 53). The final end is not directed toward any finite good and must be directed beyond the finite (MacIntyre, 2016, p. 53). For Aquinas, the virtues are qualities that allow us to overcome inward and outward evils to pursue our final end by achieving individual goods and common goods. MacIntyre conceptualizes *common goods* as goods that are achieved and enjoyed by being a member of a particular role and participating in a relevant type of activity (2016 p. 169). Aquinas argues that individuals cannot achieve their individual goods except through achieving common goods, so moral agents cannot act without also acting as political or social agents (MacIntyre, 2016, p. 182). MacIntyre clarifies that this does not mean that the goods for an individual must be subordinated to the goods of the community or vice versa. Rather, for individuals to define and pursue their individual goods in concrete terms, they must recognize their community's common goods as their own (MacIntyre, 1999, p. 109). As an agent pursues common goods, they cultivate the virtues and learn what it means to flourish.

In his neo-Aristotelian developments of flourishing, MacIntyre borrows from the insights of Sophocles and Marx to conceptualize the conflicts and social conditions of a political community. For Sophocles, conflicts which individual agents encounter affect the fate of individuals and the community, as individuals play a role in representing their community (MacIntyre, 2007, p. 208). MacIntyre explains that although the Sophoclean self remains accountable to its community, it can also transcend the limitations of its social roles and put them into question (2007 p. 145). MacIntyre argues that the Sophoclean insight missed by Aristotle is that the conflict of human life can also be one important source of learning about human flourishing; it is through conflict—and sometimes only through conflict—that we can learn what our ends and goods are (2007 p. 163-164). Drawing from Marx, MacIntyre also highlights the importance of *sociological self-knowledge* to know who you are and those around you (MacIntyre, 2016, p. 211). Sociological self-knowledge involves an awareness of the nature of the roles and relationships in which you are involved, the shared assumptions of those you interact with; and what in these roles, relationships, and assumptions can obstruct rational agency (MacIntyre, 2016, p. 213). This knowledge also includes an understanding of what possibilities there are to act in a way that could transform these roles, relationships, and assumptions (MacIntyre, 2016, p. 213). Additionally, drawing on Marx, MacIntyre emphasizes that agents need to understand how common goods are affected by the structures through which money and power are distributed in their society if they are to flourish as members of a political community (2016 p. 211). Thus, the Sophoclean self and Marxian sociological self-knowledge enable agents to pursue common goods while also putting their social order into question to refine its understanding and pursuit of human flourishing.

4.2. Neo-Aristotelian Practical Rationality and the Quantified Self

MacIntyre's neo-Aristotelianism scheme of flourishing requires that agents properly exercise their practical rationality so as to cultivate the virtues and pursue individual and common goods throughout their lives. The conclusion of a piece of practical reasoning by an agent is an action; the "Aristotelian claim is that reasoning can issue in action" (MacIntyre, 2016, p. 189). For an action to be a conclusion of practical reasoning, it must be consistent with the premises asserted in that reasoning, in which an overriding good in an agent's situation is identified, and aimed at, in the act (MacIntyre, 2016, p. 189). MacIntyre highlights that whether agents do or do not act toward the goods depends upon the quality of their practical reasoning, and "the quality of that reasoning depends in turn on how far they are able to distinguish genuine from merely apparent goods" (2016 p. 190). It is an agent's virtues and vices that determine their ability to differentiate between true and false goods in this way. The virtues are those qualities that enable agents to identify what goods are at stake in a particular situation, the relative importance of those goods to the situation, and how that particular agent must act for the sake of the best good (MacIntyre, 2016, p. 190). Thus, MacIntyre's neo-Aristotelian view contends that being a good practical reasoner is closely related to being a virtuous human being (2016p. 190).

How does the pursuit of the quantified self affect agents' ability to distinguish between genuine and merely apparent goods when exercising practical reasoning? It has been argued that the data models at the core of the quantified self do not reveal a holistically truthful account of agents' moral lives, but instead only provide a working model for predicting certain quantifiable phenomena under limited conditions. The informational reductionism of the quantified self also neglects how the intricacies of moral reality cannot be wholly represented in data models. Consequently, the things that an agents' data profile would reveal as "goods" would also be problematically oversimplified from their complex reality and detached from the moral development of the agent. The moral situations in which agents practically reason would be necessarily reduced by the datafication of the quantified self.

If an agent's data profile revealed that they lack sufficient social contact in their lives, this revelation would be the product of a calculation of certain variables (to which multiple valuations could be given) that indicate a deficiency in the agents' trait of "social contact." If the agent were to conclude that they must seek more social contact as soon as possible, as their data profile suggested, they might fail to notice salient details about their situation. For example, the data profile might not suggest how increased social contact should be best pursued. The algorithm would likely recommend social contact in general (including phone

calls, video calls, and text messages), when what the agent really needs for their moral development is a particular kind of social contact—such as face-to-face interaction with individuals who care about the good of the agent and who are willing to challenge the agent on their thoughts and desires. Moreover, the data profile might not recognize the agents' deeper needs; increased social contact may only provide a surface-level solution for the agent's moral development. It might be that the agent has had difficulty engaging in sustained social contact because they feel inadequate around others, and that inadequacy might stem from insecurities that have the vice of improper pride at their root. The agent would need important details such as these, omitted by the data model's recommendations, to truly attain self-knowledge toward their moral cultivation. Hence, the “datafied” reductionism of the quantified self would omit salient moral facts about an agent's character in a way that would hinder their ability to practically reason well.

As mentioned in the contrast between the quantified and cultivated self, the former tracks an increasing number of personal variables without a clear sense of how these variables are to be rank-ordered in terms of moral importance. As a result, the ability to differentiate between true and false goods during practical reasoning would also be muddled by the plethora of variables and the lack of a clear standard by which to judge the relative moral importance of each for the agent's moral cultivation. Tracking variables including the agent's screen time, viewing habits, run times, nutritional info, spending habits, heart rate, sleep cycles, etc., would do little to help the moral development of the agent unless there was a detailed understanding of how these variables play into the agent's moral life, and when one variable should be given more attention than the others given the agent's situation. The lack of a moral aim toward self-cultivation makes the numerous tracked variables of the quantified self a liability for quality practical reasoning.

Since moral concepts such as “happiness” and “well-being” are reduced to a subjective sense in the quantified self, agents would be unable to practically reason from an accurate understanding of their moral state. For instance, an agent's data profile might indicate that they rate highly on the “well-being” category due to a subjective sense of contentment influenced by extraneous factors such as good weather, bodily health, and energy levels. The agent might then assume that their moral state is in a good condition even though they might have serious character defects and hidden vices that have not risen to the agent's awareness due to the external factors causing a positive subjective sense of well-being. The quantified self's datafication could confuse moral excellence with subjective contentment and thereby hinder the agent's self-knowledge about the state of their present moral character. The absence of substantive moral guidance in rank-ordering personal variables and the subjective sense

given to moral categories would thus further damage the proper exercise of an agent's practical reasoning.

MacIntyre also identifies two kinds of premises characteristic of neo-Aristotelian practical reasoning. The first kind identifies the particular good to be achieved or particular evil to be averted by acting here and now toward a particular end (MacIntyre, 2016). The second kind are premises about facts pertaining to how to act to achieve that end, including the constraints under which the agent must act if they are to not potentially cause bad unintended consequences (MacIntyre, 2016, p. 190). MacIntyre highlights the different formal structure between practical reasoning and theoretical reasoning. In theoretical reasoning, we can add true premises to a sound argument and it will remain sound, but with practical reasoning, additional true premises may no longer make the argument sound (MacIntyre, 2016, p. 191). MacIntyre gives the illustration that if I need to catch a plane, and it is true that "I can catch it if I leave now," the additional true premise that "the flight is cancelled" now makes the conclusion "I should leave now" void (2016 p. 191). Practical reasoning thus demands that agents ask what other goods are at stake in a situation besides those that are immediately noticeable. Agents must consider what are *all* the relevant goods to be considered in a situation when exercising practical reasoning. For MacIntyre, sound practical reasoning requires that an agent "think or have thought more widely about the range of individual and common goods that it is open to him to achieve" (2016 p. 191). To fail to consider the range of relevant goods is a failure in the exercise of the virtues which enable agents to imagine a range of achievable goods in diverse situations (MacIntyre, 2016, p. 191).

How does the quantified self hinder the requirement of practical reasoning to consider all the relevant possible goods in a situation? Since the quantified self movement seeks to reduce moral reality and the virtuous character of agents into machine-readable data points, reducing agents' apprehension of moral reality is the *modus operandi* of its data models. When surveillance technologies record and analyze personal data with AI algorithms, they must prepare and interpret that data into quantifiable variables to input into individual data profiles. In this process, the reductionism produced by datafication (caused by both the algorithm and the human decisions in the data science process) makes it such that agents are unable to see all the relevant goods at stake in a situation if they are only—or even primarily—informed by data models. If agents believe in the vision of human life cast by the quantified self, they will only be able to imagine a range of goods in a situation that accords with the algorithmically generated variables, and thus, will inevitably be blind to many relevant goods in situations throughout their lives as they try to practically reason. The quantified self thereby hinders practical rationality by limiting the scope of salient goods that agents can "see," and perhaps even imagine, through its quest for increasing datafication and corresponding informational reductionism.

Conclusion

Both claims made by supporters of the quantified self are thus shown to be dubious. The supposed enhanced objectivity and truthfulness of the data profiles of the quantified self are found to be, in fact, reductions of complex moral reality that are permeated with subjective human decisions throughout their construction and application. The alleged contribution to human flourishing by the quantified self via improved self-knowledge turns out to be misleading when assessed by a virtue ethics approach that understands self-examination as a moral practice for moral self-cultivation. By applying the concept and practice of neo-Aristotelian practical reasoning to the quantified self, agents are found to be hindered in their perception of a situation's possible moral goods and are apt to misunderstand their own moral character when relying on feedback from data models. Additionally, inhibited practical rationality may also prevent agents from exercising the sociological self-knowledge to adequately put their own social order into question. The implication is that the adoption of the pursuit and worldview of the quantified self in a social order might impede agents inhabiting that order from recognizing and questioning the assumptions, subjective human components, and reductionism of the quantified self—a noteworthy consideration for those interested in forming a more truthful and morally wise society.

References

- Ajunwa, I. & Schlund, R. (2020). Algorithms and the Social Organization of Work. In M. D. Dubber, F. Pasquale & S. Das, *The Oxford Handbook of Ethics of AI* (pp. 805-822). Oxford: Oxford University Press.
- Azati. (2021, June 3). *How Artificial Intelligence (AI) Is Used in Targeted Marketing*. Retrieved December 15, 2020, from Azati.ai: <https://azati.ai/artificial-intelligence-targeted-marketing/>.
- Coeckelbergh, M. (2020). *AI Ethics*. Cambridge, MA: The MIT Press.
- Francis, L. P., & Francis, J. G. (2017). *Privacy: What Everyone Needs to Know*. New York: Oxford University Press.
- Hildebrandt, M. (2013). Slaves to Big Data. Or Are We? *9th Annual Conference on Internet, Law & Politics*, 1–21.
- MacIntyre, A. (1999). *Dependent Rational Animals: Why Human Beings Need the Virtues*. Chicago: Open Court.
- MacIntyre, A. (2007). *After Virtue* (3rd ed). Notre Dame: University of Notre Dame Press.
- MacIntyre, A. (2016). *Ethics in the Conflicts of Modernity: An Essay on Desire Practical Reasoning and Narrative*. Cambridge: Cambridge University Press.
- Mayer-Schonberger, V., & Cukier, K. (2014). *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. London: Eamon Dolan/Mariner Books.
- Moore, P. V. (2017). *The Quantified Self in Precarity: Work, Technology and What Counts*. New York: Routledge.
- Oakeshott, M. (2001). *The Voice of Liberal Learning*. Carmel: Liberty Fund.
- Quantified Self. (2022). *What is Quantified Self?* From Quantified Self: <https://quantifiedself.com/about/what-is-quantified-self/>.
- Quantified Self Institute. (2016). *About*. From Quantified Self Institute: <http://qsinstitute.com/about/>.
- Spaulding, N. W. (2020). Is Human Judgment Necessary?: Artificial Intelligence, Algorithmic Governance, and the Law. In M. D. Dubber, F. Pasquale & S. Das, *The Oxford Handbook of Ethics of AI* (pp. 376-402). Oxford: Oxford University Press.
- Stieglitz, S., Mirbabaie, M., Ross, B., & Neuberger, C. (2018). Social media analytics—Challenges in topic discovery, data collection, and. *International Journal of Information Management*, 156–168.
- Vallor, S. (2016). *Technology and the Virtues: A philosophical guide to a future worth wanting*. New York: Oxford University Press.

Wolf, G. (2018). *Article 27*. From Quantified Self: <https://quantifiedself.com/wp-content/uploads/2019/02/Article27-Pamphlet.pdf>.

Wolf, G. I. & De Groot, M. (2020). A Conceptual Framework for Personal Science. *Frontiers in Computer Science*, 2(21), 1–5.

Zuboff, S. (2020). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.