Digital Enlightenment Yearbook 2014
K. O'Hara et al. (Eds.)
IOS Press, 2014

doi: 10.3233/978-1-61499-450-3-119

Personal Data Empowerment and the Ideal Observer

Reuben BINNS [1]
Faculty of Business and Law/Department of Electronics and Computer Science, University of Southampton, SO17 1BJ, UK

Abstract. This chapter explores the potential for market-driven personal data empowerment. It begins by discussing the results of a recent study into the economic and business case for personal information management services (PIMS), providing an overview of the variety of services and some key market developments. The following sections consider the problems that PIMS might be able solve, and the extent to which their solutions could be genuinely empowering, by reference to some key aspects of Enlightenment thought, starting with Adam Smith's notion of enlightened self-interest, which later developed into the concept of *homo economicus* at the heart of classical economics. I argue that this ideal became ever less realistic as the economic system it theoretically justified became ever more complex, with the result that in many markets, businesses profit from the fact that individuals are ill-equipped to deal with said complexity. I propose that the seeds of a solution can be found in the notion of an 'ideal observer', also originating in Enlightenment thought; a hypothetical ideal version of an individual, with perfect knowledge and rationality, whose perspective and insight can steer its ordinary, fallible human counterpart towards better decision and action. In so far as PIMS can help individuals approximate this abstract agent, they offer the possibility of genuine enlightenment and empowerment, fit for the complexity of modern consumer markets.

Keywords. personal data, personal information management services, personal data stores, privacy, adam smith, economics, enlightenment, empowerment, cognitive bias, paternalism

Introduction

Personal data is widely touted as a new economic asset. For better or worse, it is transforming whole industries, from advertising to health care, and the long-term consequences of this change are still unfolding. The potential for personal data to provide new insights and efficiencies, whilst simultaneously undermining individual autonomy, privacy, equality and other social and democratic principles, makes its use and abuse a key determinant of our future individual and societal welfare.

Alongside this familiar narrative – which is usually couched in terms of 'the good or bad things that institutions and corporations might do with your data' – there is an emerging alternative model, in which individuals become empowered by using their own

[1]Reuben Binns, Faculty of Business and Law/Department of Electronics and Computer Science, University of Southampton, SO17 1BJ, UK. E-mail: rb5g11@soton.ac.uk.

data for their own purposes. This model is manifested in the rise of Personal Information Management Services (PIMS), which help individuals gather, manage and use information – both about themselves and the world 'out there' – to improve their decisions and manage their lives better. Some believe that this model represents a new opportunity in its own right – a market for data-driven services, created for and controlled by individuals, centred around their needs and desires.

This phenomenon has been analysed in a series of research reports by industry analysts Ctrl-Shift, the latest of which assesses the potential economic and business impacts of PIMS [10]. It estimates the total value of a fully mature PIMS market at £16.5 billion in the UK. During 2013, new PIMS launched at a rate of one per week. There are a wide variety of offerings, with different funding models, value propositions and motivations. There is also evidence of growing investor interest in PIMS.[2]

The following section (2) summarises some of the major findings and themes of this research, giving an overview of the emerging PIMS market. Subsequent sections frame these developments in terms of a theory of personal data empowerment. Section 3 identifies one of the primary problems that PIMS could solve – namely, the inability of individuals to adequately deal with the complexities of modern markets and institutions – and traces the roots of this problem back to an increasingly unworkable interpretation of the Enlightenment ideal of autonomous rational individuals. Section 4 presents another Enlightenment concept – that of the 'ideal observer' – and argues that by striving to embody it, PIMS can overcome these problems and deliver genuine personal data-driven empowerment in the digital age.

1. Personal Information Management Services

1.1. What are PIMS?

Personal information management services can be seen as the latest phase in a longer trend of digital disruption, beginning with the advent of cost-effective information processing and storage. Just as home computing and the web allowed individuals to do things which were previously the preserve of large organisations (such as using spreadsheets or searching vast troves of data), so PIMS enable individual equivalents for organisational uses of personal data. Principally, they allow individuals to gather, store, manage, use, and share their personal data under their own control.

These basic functions of personal data management are provided by *personal data stores*,[3] which act as a locus of control and point of integration of personal data around the individual. A PDS is more than just a personal file storage system for documents and media, such as currently popular cloud-based storage services.[4] Rather, it is a platform for integrating many kinds of personal data from a range of external sources (such as

[2] Of the ∼300 PIMS identified by Ctrl-Shift during the report, public data on venture capital investments could be found for 99 of them on CrunchBase.com, a directory of technology companies. A rudimentary analysis for the purposes of this chapter found that 49 of those have some level of venture capital investment, a total amount of $457 million.

[3] Personal Data Stores (PDS) may also be referred to as personal data *vaults, lockers, and clouds*.

[4] Examples of currently popular cloud storage services in the US and EU market include DropBox and Google Drive.

personal devices and bank accounts), sorting, normalising and categorising it, and selectively releasing it to external locations.

The data they contain could be attributes – such as name, age, or location – but also more dynamic feeds, such as fitness data or recent purchases. As such, most PDSes run fully-fledged databases, and have application programming interfaces (API's) which allow them to communicate over the Internet with other programs. In these respects they are dynamic applications, able to respond to events and trigger actions. While most of the PDS systems on the market run in the 'cloud' (i.e. on servers managed by the PDS provider), some allow individuals to run their own PDS on their own device or choice of host. Others adopt a hybrid model which maintains both individual control and the benefits of the cloud by giving users exclusive access to the means of encrypting and decrypting their data. In this way, the PDS provider relinquishes the capacity to surveil its users despite hosting their data.

Some PDSs are intended to be general-purpose and extensible to all manner of use cases, while others are more specific, dealing with particular types of data or offering certain kinds of functionality. Some users may favour a single, centralised PDS, while others might prefer to spread their data management across different stores; both strategies have their strengths and neither is more paradigmatic of the PDS approach.

By providing the underlying infrastructure for managing personal data, a PDS can in turn enable other services. One important category is *decision support services* (DSS) which help individuals make more appropriate purchases and lifestyle choices. An early form of decision support that has already gone 'mainstream' is the price comparison website, which aggregates product pricing information in one place so that users can make easy comparisons. Whilst such services are useful, to be really effective, decision support needs to be contextually relevant. The most powerful DSSs combine information about the individual – such as their circumstances, behaviour, values and preferences – with information about the world 'out there', such as pricing structures, product features, peer reviews, impartial ratings and environmental conditions, crunching these two categories of data to arrive at a personal recommendation.

An exemplar DSS is BillMonitor, which analyses an individual's mobile phone usage data against a comprehensive database of telco providers' tariffs to work out their optimal plan. This is an example of a calculation that is simply too complex for even motivated and skilled consumers to carry out for themselves, as it must take into account their own historical usage data over a period of months, and compare that to the thousands of different tariff permutations on offer. Services like this allow consumers to optimise their decisions to a previously impossible degree.

When making purchasing decisions, the average consumer is often racked with uncertainty over what constitutes quality in a given product category – which features are important and which are unnecessary? A DSS like *JustBuyThisOne* can provide clarity. It sifts through complex technical information to answer these questions, delivering a single recommendation for each product category. Some decision-support services tell you not only what to buy, but when; forecasting whether prices will rise or fall and providing alerts when they drop below a certain threshold.[5]

Ethical purchasing is another growing area for DSSs. Services like GoodGuide and EthicalConsumer allow users to build a personal profile of their ethical concerns on a

[5]For instance, Slice provides product price drop alerts, while Kayak, a flight ticket intermediary, makes price change forecasts with a confidence interval based on historical data.

sliding scale, from broad categories like 'Environment', 'People' or 'Animals', to specific factors like carbon emissions, responsible supply chain management, workers' rights, or factory farming. These profiles drive what EthicalConsumer describes as a 'personal ethical rating system, based on detailed research of over 40,000 companies, brands and products', to help users find the products that best meet their unique combination of ethical concerns. Similarly, a number of food-related applications crunch data on common ingredients to allow users to select products that match their complex dietary requirements, such as vegan, halal, kosher, gluten-free, 'paleolithic', or even some combination thereof.[6]

The third main category of PIMS is *life management services* (LMS). They help individuals manage various aspects of their lives better by automating the planning, scheduling, organising, coordinating, integrating, and administering of tasks and processes. LMSes often focus on a particular sphere such as finances, home, or travel, working across organisational silos to reduce confusion, time-wasting and hassle for their users. For instance, personal finance managers pull in data from multiple bank accounts to provide users with a single view of their finances, from which they can identify trends, track spending, and set savings goals. CardCosts, a tool provided by the UK trade body for credit cards, helps individuals navigate the complex world of credit card fees, calculating the optimal debt payment strategies for an individual given their circumstances and their bank's terms and fee structure. Services like BillGuard scan credit and debit card purchase histories to automatically identify erroneous charges and put users in touch with retailers to resolve issues.

The identification and rectification of compensable or unnecessary spending is a recurring theme in LMSs. For instance, a number of travel management apps help travellers claim refunds automatically in the event of delays or cancellations.[7] They scan travel card and email accounts for records of journeys, cross-referencing them with public transport data to identify journeys that have been delayed to a degree sufficient to make the individual eligible for a refund. The services then automatically file claims with the relevant transport provider, requiring no input from the individual. This is also part of an evolving trend towards tools which help individuals manage their purchases after money has changed hands; other examples include consumer alert services which inform consumers when flaws in the products they own are discovered, or if they have been recalled by the manufacturer.[8]

A nascent market for LMSes is independent third-party services which augment or bypass traditional customer service channels. Successful dispute resolution via a company's own channels can be an information-intensive process, requiring consumers to maintain meticulous logs of their complaints, pay attention to time periods, billing information and contract clauses, and navigate complex automated phone tree systems. A third party acting on behalf of the consumer can reduce these burdens. For instance, Resolver is a UK service which helps consumers make complaints to over 1,500 companies. It stores their complaint information in one secure place, guides them through the complaint process for the company in question, and tracks progress towards a resolution.

Generic life management software like calendars, to-do lists, and contact managers have been in use for several decades, but a new generation of services aims to crunch the data inside these tools to provide intelligent reminders, recommendations, and advice to

[6] See, for instance, Klappo and Ingredient 1.

[7] Examples include AirHelp (for air travel) and ClaimMyRefund (for public transport in London).

[8] Examples include Slice and Smarter Complaints.

their users. The common use cases for these 'personal digital assistants' are reminiscent of the applications of artificial intelligence imagined in science fiction. For instance, by accessing a user's location and calendar data, a digital assistant might send them an alert when they need to leave their current location in order to arrive on time for a meeting across town, providing information on the best route and transport options, bringing up relevant information on those attending the meeting and documents to read in preparation.[9] One of the main challenges for such applications is pulling in the right personal data from other silos; which may explain why existing consumer technology giants, who already have access to those silos, are active in this area.

Many services combine elements of life management and decision-support to provide both better decision-making and easier implementation of those decisions. For instance, Cheap Energy Club helps households find the cheapest energy provider currently available to them. Users upload their relevant information – including location, energy usage, and savings targets – to a PDS. The service then regularly scans the market, alerting the user whenever a new energy deal better matches their criteria, and helps facilitate the process of switching supplier. By automating the boring administrative tasks of searching, comparing and switching between providers, the service drastically reduces the hassle which prevents most consumers from getting the best deals in the home energy market.

A number of features distinguish PIMS from traditional providers of products and services. The major difference is that they tend to act in the capacity of an intermediary, concierge, broker or adviser on behalf of consumers. In many cases, a core part of PIMS marketing is the purported alignment of their own and their users' interests; some even emphasise how their activity runs counter to the aims of the existing brands they help consumers deal with.

They also typically help individuals manage their relationships with many brands, rather than each brand having their own tools to manage relationships with individuals and maintain data about them (in the form of Customer Relationship Management systems). They enable individuals to specify their own wants and needs, using this information to connect them to suppliers through channels the individual can control. They change the way typical consumer markets work by matching supply to demand through new platforms, driven by new sources of data.

If they go mainstream, the impact of PIMS on existing businesses is likely to be significant. On the positive side, by facilitating the sharing of volunteered, rich, quality data from individuals to organisations, they open up opportunities for new information services and more informed product development. Communication channels could also become more efficient and effective, giving organisations a better chance to communicate with individuals who have given them permission. The same mechanism could also improve risk estimation and compliance activities like fraud detection.

However, by helping their users switch suppliers more easily, PIMS could increase customer churn, with additional costs falling on suppliers. They could become the primary channels through which individuals engage with suppliers, threatening to usurp many brands' existing position as middle-men between producers and consumers. They might also subsume existing products and services into new integrated bundles, reducing existing brands' offerings to mere commodities. By giving individuals the means to manage how their personal data flows to organisations, they could upset current marketing practices

[9]This is a flagship feature of the Google Now software, as well as competitors like Donna, TempoAI, 24me and ReQall.

and other business processes which rely on easy access to personal data. Finally, they are likely to disrupt existing third parties such as advertising agencies, data brokers, credit reference agencies, customer relationship management (CRM) software vendors, loyalty scheme operators, and market research companies.

Existing PIMS have adopted a variety of business models. Most are free to their users, and charge organisations fees for (permission-based) access to users' attention and data. These fees can be either subscription or transaction-based, including commission. In some models where personal data is traded, revenues from commission fees are shared with users (see Section 1.2). In other cases, consumers simply pay for the service directly.

Ctrl-Shift modelled these revenue streams and business models using data gathered from a range of existing PIMS and traditional organisations in various industries (Ctrl-Shift 2014). On the basis of this model, the total value of this potential market is estimated at £16.5 billion in the UK (equivalent to 1.2% of the UK economy).[10] This is made up of both subscription fees paid by organisations for connecting to PIMS with the individual's permission, valued at £11.5 billion, and commission fees paid for one-off transactions valued at £5 billion. This gives an indication of the potential significance of these emerging services; for comparison, this would place the PIMS market above the automotive (0.7%) and pharmaceutical industry (0.97%) in terms of contribution to the UK economy.

1.2. Personal Data, PIMS, Empowerment and 'The Market'

There are several ways in which the concept of the market enters into a discussion of personal data empowerment. One way is the 'market' for personal data, where data itself is traded or exchanged much like a commodity.[11] Another sense (which is discussed above) is the market for personal information management services, where the 'size' of this market is the amount of revenue these new entities could gain. Finally, we might talk about the impact of PIMS on other markets, how they are structured and the way they operate. Each of these distinct references to the market have important implications for how we understand empowerment, which are worth exploring briefly.

1.2.1. The Market for Data

Some PIMS seek to empower their users through a modification of the status quo, where personal data is a good to be bought and sold by organisations. These PIMS services offer to act as brokers of personal data; the difference with the current 'wild west' data broker model is that in this case, individuals get a fair share of the profits. For instance, services like *Handshake, DataCoup* and *Knowledge Levers* offer to store their users' data, and rent it out to selected marketers with their explicit permission, returning a share of the profits to the individual.

For an industry accustomed to paying a small unit price for large volumes of data gleaned by data brokers (usually without the knowledge or consent of individuals), the benefits of this alternative model may not be immediately obvious. Why pay for access

[10] This represents the value of a fully mature PIMS market, aside from any costs associated with creating such a market.

[11] There are difficulties in classifying personal data as a genuine commodity. It is infrequently paid for (more often traded for services) and therefore its 'price', where it is paid for, is unlikely to be a function of an established market. Neither is it fully 'fungible'. Nevertheless, commoditisation is a spectrum and in some cases personal data is traded much like a commodity.

to a handful of personal data (even if volunteered and of high quality), when you could cheaply buy targeted ad impressions by the thousand? One answer is the potential to offer new, superior quality data through a voluntary system. Ensuring that this higher-quality data remains scarce enough to be an attractive option for existing buyers such as direct marketing agencies remains a challenge. For some, the solution is technical, by limiting access to data through security measures, access control and sticky policies [11]. For others, the answer lies in contract or property law, which can create artificial scarcity, in a strategy similar to that employed (with varying degrees of success) in the copyright industry [12].

Whatever the solution, presenting empowerment solely in terms of giving individuals control over how their data is sold is problematic. As stated in the introduction to the last edition of this volume, 'personal data management should not simply endow the data subject with the ability to sell to the highest bidder' [13]. These services may have some immediate appeal in allowing consumers to play the market, game the system, and 'get their own back' on what many regard as a nefarious data broker industry. As suggested in the title of a recent book about the monetisation of personal data, '*Das Kapital, Bin Ich*' ('I am capital'); consumers – or more accurately, their *attention* – are the product being sold [14]. The allusion to Marx might suggest a possible solution in communal ownership of the means of (personal data) production, or a data-subject cooperative. But this already cedes too much to the status quo; the opportunities for genuine empowerment within existing personal data markets may be limited.[12]

1.2.2. The Market for PIMS

There is a difference between a market for personal data itself, based on the monetisation of that data, and a market for PIMS. PIMS may in some cases (as we have seen above) participate in the market *for* personal data. But this is not their primary purpose. The PIMS market isn't a market for personal data, but for *services* which use that data in new ways to serve a broad range of individual needs, desires, values and goals. Just how such services emerge and whether they achieve critical mass remains to be seen. The figures referenced above give an indication of the potential value of a mature PIMS market, but they do not predict or explain how that market might evolve. That will depend heavily on the reaction of regulators, the effect on industry bottom-lines, technological development and public trust.

1.2.3. PIMS as a Market-changer

While it reveals a wider picture than the personal data market alone, focusing on the market for PIMS doesn't capture their true import either. PIMS are significant because of the ways that they could change how markets operate in general. These include equipping consumers with tools to engage with sellers through channels they control, on their own terms; providing the ability to gather, store, curate and share rich, verified personal data

[12]It is worth considering here the differences and relationships between empowerment and equality. As economic units in these markets, individuals don't start out with an equal distribution of capital. The buying habits of the spendthrift rich may be worth more to marketers than those of the prudent poor. In this sense, the benefits of empowerment arising from individual-controlled data brokerage service (and, perhaps, PIMS more generally) are likely to be unequally distributed amongst the population, possibly strengthening existing inequalities or creating new ones. So it may be that PIMS empower unequally; however, it is neither necessary nor sufficient that PIMS be *equitable* in order for them to be *empowering*.

with selected organisations, thus enabling powerful new information services; and giving consumers new leverage and bargaining power over suppliers of goods and services who have to work harder to gain their trust and custom.

To be truly empowering, PIMS cannot simply be a good investment or drive incremental improvements in consumer convenience. The ultimate mark of success for PIMS-driven empowerment is the extent to which they succeed in changing markets and increasing the power of individuals who operate within them. To that end, we need a theoretical framework to understand PIMS' empowering potential and the nature of the problem they are an answer to. I propose that such a framework can be constructed from concepts originating in Enlightenment thought (in the spirit of this publication), in particular in the work of the philosopher and proto-economist Adam Smith.

2. The Origins of Disempowerment

The discussion thus far has assumed that personal data empowerment is a necessary or desirable end. But in what senses are individuals not already empowered, or disempowered? What is the problem to which PIMS are the proposed solution? In many public debates, the presumed problem is referred to by the catch-all term of *privacy* (or the lack of it). Other worthy candidates include concerns about 'digital serfdom', or the spectre of big data and its gradual erasure of civil rights [15]. On the other hand, the stated goals of PIMS themselves are often mundane – such as saving time on administrative tasks or finding cheaper electricity – with more fundamental values such as privacy and civil rights taking a backseat.

Whether lofty or mundane, these are all worthwhile goals. But there is a more pervasive problem for which PIMS are a possible solution; that is, the general inability of individuals to make informed, meaningful and optimal choices, and to efficiently implement those choices, given the complexity of the modern world. PIMS have the potential to fix this, providing individuals with the capacity for a level of agency greater than ever previously possible.[13] To understand both where this complexity came from, and how PIMS can be a solution to it, I turn to the intellectual origins of the modern marketplace. In presenting this ad hoc selection of intellectual history I make no claim to rigorous historical analysis or exegesis, but simply offer what I hope are enlightening theoretical constructs.

The theory underpinning our current economic model arguably has its birthplace in the Enlightenment, in particular in the work of Adam Smith.[14] In *The Wealth of Nations* [16], Smith famously argued that in pursuing their own self-interest in the marketplace, individuals inadvertently promote each others' respective utility (in this context, 'utility' is often characterised, after a little philosophical hand-wringing, in terms of *happiness* or *satisfaction*). When individuals pursue their own interests, according to Smith, the

[13] Of course, levels of agency vary between and within societies and are rarely equally distributed. Any benefits of PIMS are likely to be limited initially to their target user base which is, for the most part, relatively affluent individuals in economically developed countries. The following discussion should be read with this caveat in mind, although it is the hope of the author that the scope of PIMS broadens to encompass (or indeed, specifically target) the relatively disadvantaged.

[14] Indeed, the economic status of the 'developed' world has been attributed by some to the Enlightenment itself [2, 5].

'invisible hand' of the market ensures that their efforts also benefit others, often more effectively than attempts to promote the common good directly.

> It is not from the benevolence of *the butcher*, *the brewer*, or the *baker* that we expect our dinner, but from their regard to their own *interest.*

Smith's thesis does not imply that human nature is, as a matter of fact, inherently selfish, nor is it a normative endorsement of egoism (indeed, Smith goes to great lengths to refute both positions in his earlier work, the *Theory of Moral Sentiments*). However, Smith's insights had a profound influence on the foundations of classical economics, with the *Wealth of Nations* later credited as providing 'the most important substantive proposition in all of economics' [17]. As a result, Smith became associated with a set of fundamental assumptions about human nature which were imputed to him, despite being at odds with the subtleties of his own views. At the heart of this classical model is the notion that individuals are best placed to choose rationally from the multitude of options that which will bring them most happiness, and furthermore that individuals do indeed generally behave in this way.

This idealised individual, more recently given the cod-Latin name *homo economicus*, fed into Jevons' theory of marginal utility in the 1870s, thus giving us economics as we know it today. Crudely summarised, it says that we all buy and sell so as to increase our individual utility, such that demand and supply tend to equilibrate and the prices we accept are a good proxy for the marginal utility resulting from a sale or purchase.

A full discussion of the theoretical merits and failings of *homo economicus* is beyond our present scope, but the following points are particularly germane to our discussion. First note that, according to this theory, buyers and sellers are essentially the same kind of agent, capable of applying the same resources and capabilities to their commercial decisions. But in contemporary consumer markets,[15] buyers are generally individuals while sellers are more usually large organisations, and each comes to the market with radically different capabilities, tools and levels of information.

Some complexity in consumer-facing markets is deliberately engineered by the companies operating within them, and even more complexity arises unintentionally as a by-product of their activity and a globalised economy. One example is the complexity of processing pricing information; in 2011, there were 8.5 million different mobile phone tariff combinations available to consumers in the UK alone.[16] The administrative complexity of many consumer tasks like claiming refunds or switching suppliers means that the fundamental rights of the free market are rarely exercised. And prudence is not the only consumer value thwarted by complexity; ethical and environmental concerns are similarly hard to translate into action. There is a large, dense and often contradictory body of literature and media stories exhorting the ecological benefits and harms of various products, leaving most consumers flummoxed and disengaged. Assessing the rights and responsibilities associated with a good or service is another; the digital music service iTunes explains its terms of use in a 56-page document,[17] while online payment processor PayPal's user agreement is longer than Shakespeare's Hamlet.[18]

[15] Of course individuals face similar challenges outside of their 'consumer' identity, in their capacity as citizens or employees for example. What follows primarily concerns individuals *qua* consumers.

[16] According to http://www.billmonitor.com

[17] See http://edition.cnn.com/2011/TECH/web/05/06/itunes.terms/

[18] From http://conversation.which.co.uk/technoloigy/length-of-website-terms-and-conditions/

In each case, an ideal individual would consider all the information relevant to a decision, both about their own unique circumstances, needs, and preferences, and the world 'out there'. But given all this complexity, even the most diligent person is unlikely to have the time and inclination to take in and process all the relevant factors when making a purchase. Even if they were so inclined, optimal decisions usually require more processing power than the human brain is equipped with. Rather than maximise our utility, at *best* our behaviour yields merely sufficiently satisfying rewards given our cognitive limitations. In Herbert Simon's words, we exhibit *satisficing* behaviour [18]. And as later work in behavioural economics has revealed, we may be systematically biased towards various forms of irrationality, such that even satisficing behaviour is an unachievable ideal.[19]

The consequence of all this is that in many consumer-facing markets, profits are made by exploiting these impairments and biases, and the consumer inertia, ignorance, and impotence that they create.[20] The result is that, in practice, classical Economics' vision of the rational, autonomous, utility-maximising agent is out of reach. Its legacy is an unfortunate combination; complex consumer markets and an ideal of consumer agency that is increasingly unworkable in the face of that complexity.

I propose that this is the fundamental problem to which PIMS could be a general solution. They provide an additional layer in between consumers and traditional suppliers of goods and services, acting on behalf of the former to engage with the latter. They bring value to individuals by gathering data and analysing it in order to cut through the complexity of confusing, difficult and boring consumer-related tasks.

3. The Ideal Observer

3.1. Maintaining Agency

If the basic premise of PIMS is to support, augment or otherwise enhance individuals choices in complex markets, some important questions immediately arise. What sort of choices should they help individuals make? What sort of values should they embody? And can this be done in a way that empowers rather than undermines their agency?

If the primary purpose of personal information management were to simply correct for our failure to act like *homo economicus*, then their most obvious goal should be to make consumer purchases more like those of *homo economicus*. This would undoubtedly save consumers money in markets where competition and regulation has failed to drive down prices, and indeed this is the model followed by many of the PIMS featured above. Alternatively, they might move on from this ideal, instead embodying something broader; perhaps reflecting our altruistic, aesthetic and political values. For instance, a service which helps individuals manage their home energy might factor in their environmental as well as prudential concerns.

Whatever ideals PIMS help us embody, there is a danger that the decisions and actions that arise from them cease to be attributable to the individuals they act on behalf of. Users may begin to feel the decisions and actions the service undertakes are not implementations of their own will but rather impositions from an outside force. If individual

[19] See [4] for an overview of research in this area.

[20] This can be true both in the usual case where the individual is the buyer of a product, and as with many digital platforms where individuals themselves are 'the product' being sold to advertisers.

agency becomes disassociated from the service it is enacted through, then in what sense is the service genuinely empowering? On this view, rather than restoring individual agency, PIMS represent the final nail in its coffin.

3.2. 'That Great Arbiter of His Conduct'

A promising answer to this problem of agency lies in the concept of an 'ideal observer', which also originates in Enlightenment thought. The general idea is that it is possible to consider not just how individuals actually do choose and act; but how they would choose and act given greater levels of knowledge and the capacity for rational deliberation. Despite being hypothetical, such choices might still be objectively and inextricably linked to the individual as a function of thier existing character, circumstances, values and preferences, as opposed to choices made by an outside agent with different values.

An early version of the ideal observer theory can be found in Smith's own work, as well as in that of his compatriot David Hume.[21] Seventeen years before *Wealth of Nations*, in his *Theory of Moral Sentiments*, Smith identified the gap between our actual cognitive abilities and those of a hypothetical rational agent (in this, Smith arguably pre-empted the work of Simon and the behavioural economists [19][22]). He introduced a hypothetical agent – the 'impartial spectator' – as a way to understand how we sometimes overcome our bias for short-term gratification. The impartial spectator is a version of ourselves capable of:

> self-denial, of self-government, of that command of the passions which subjects all the movements of our nature to what our own dignity and honour, and the propriety of our own conduct, require (Smith, 1759 [1981], I, i,v, 26).

For Smith, a person arrives at wisdom by aligning his desires, intentions and feelings with those of 'that great arbiter of his conduct' – his individual version of the impartial spectator.

The notion of an ideal observer was further developed in 20th century analytic philosophy in the search for an objective and non-paternalistic notion of what is 'good for' an individual, or what is genuinely in their best interests [20–22]. Peter Railton argues that what is good for a person is not necessarily the object of their present desires, nor what some external ethical code says they should want. Rather, the ultimate measure of what is best for someone is what their ideal self, that is, one 'fully and vividly informed about himself and his circumstances, and entirely free of cognitive error or lapses of instrumental rationality' would want [22]. Similarly, for Bernard Williams, claims about what an individual *should do* are equivalent to the result of informed deliberation starting from their current motivations; a logical deliberative route must be traceable from an individual's current motivations to her putative 'ideal' choices [21].

This way of thinking allows us to accept that an individual is capable of misjudging what is in their own interest, whilst maintaining that an accurate assessment of their genuine best interests must proceed from their existing values, personality and motivations. The ideal observer is not an external hector laying down universal moral laws, but instead recognises the nuance and detail of a particular situation and the unique attributes of the

[21] It is testament to the heterogeneity of the Enlightenment tradition, and the capacity of subsequent generations to selectively develop it, that we find within it both the origins of the problem and the seeds of its solution.

[22] In light of this, Smith has even been interpreted by some as a proto-behavioural economist [8].

individual within it. Ignore this, and discussions about an individual's 'best interests' risks treating them as a mere channel between the input of 'the utility network which the projects of others have in part determined' and 'an output of optimific decision' [21]. In Williams' words, we neglect

> the extent to which *his* projects and *his* decisions have to be seen as the actions and decisions which flow from the projects and attitudes with which he is most closely identified. It is thus, in the most literal sense, an attack on his integrity.

This an important consideration for designing PIMS which seek to empower individuals to serve their best interests without undermining their integrity and autonomy. On the one hand, we know that individuals' existing choices should not be taken at face value; this is the conclusion of the critiques of *homo economicus* articulated by Simon and the behavioural economists, as well as from the examples of consumer market complexity given above. But this doesn't mean that PIMS need to be paternalistic, ignoring their users' existing motivations and imposing new choices and behaviours upon them externally. Such services would be at odds with the notions of individual empowerment, autonomy and integrity. Instead, PIMS design should take the individual's existing motivations, values and preferences as the starting point for decision support and implementation.

The formula for taking an individual's existing motivations and arriving at decision and action need not be direct or rigid. Rudimentary PIMS might take information about the individual and the world and derive certain automated decisions and actions in a formulaic way, but more sophisticated services couldn facilitate a less direct, more iterative approach, whereby the individual has multiple opportunities for input into the algorithmic processes. This could help users to understand their behaviour and needs, and guide them to resolve potentially conflicting preferences, rather than simply computing the 'correct' result from their current motivational state.

One way to conceive the resulting relationship between an individual and their chosen PIMS is in terms of 'parameter-driven delegation'.[23] That is, PIMS take responsibility for decision-making and implementation on behalf of their users, but within parameters their users define and shape at various stages. This has the potential to eliminate some of the most debilitating instances of consumer inertia (e.g. the hassle of changing bank accounts or energy suppliers) which mean that most consumers stay with the same provider for life, often to their long-term detriment. For PIMS, there is a business opportunity in helping individuals to overcome this kind of inertia and exert power in a way that follows from their prudential, practical and ethical concerns.

3.3. PIMS and Paternalism

The emulation of an ideal observer also suggests how PIMS might be able to sidestep the accusation of paternalism which has plagued proponents of the 'nudge' philosophy (principally Cass Sunstein and Richard Thaler [23]). Nudge is an approach to public policy which advocates changing the 'choice architecture' – the way in which a range of options are presented to individuals – so as to 'nudge' individuals to make socially optimal choices. Critics argue that argue that Sunstein and Thaler's notion of 'libertarian paternalism' is an oxymoron. The authors counter that paternalism is unavoidable in situations where individual choices are unconsidered responses to their choice architecture [24].

[23] A term coined by Alan Mitchell in *Next Generation Intermediaries* [7].

PIMS might offer a way through this impasse by enabling users to set their own nudges, asserting more control over the shape of certain environments (especially digital environments), and thus control over their routine behaviour.[24] In an age where individuals are subliminally influenced, on the one hand, by carefully targeted advertisements, and 'nudged' on the other by benevolent regulators, the ability to 'nudge yourself' – that is, to set one's own choice architecture in accordance with the behaviour one wants to promote in oneself – becomes an essential capacity for agency.[25]

If they offer this functionality, PIMS might end up as a constraint on choice, for instance, by screening out certain products which fail to meet certain requirements. But this needn't undermine autonomy if those constraints are set and endorsed by the individual; self-imposed constraints on choice can promote autonomy rather than undermine it. The classical example of this comes from the story of Odysseus, who shackled himself to his ship's mast so as to hear the sirens' call without the risk of giving in to their temptations. Similarly, users of decision-support services might deliberately constrain their ability to choose, in the service of some higher-order personal goal.

3.4. Designing Transparency and Trust

In attempting to embody the ideal observer, PIMS will face significant engineering and design challenges, including getting the right data, taking into account the right values, measuring them appropriately, and crunching them to arrive at the right decisions, involving calculations that are becoming 'ever less tractable' [13]. PIMS will have to demonstrate that all the relevant individual interests – from privacy to prudence, practicality and fun – can be computed and factored into their algorithms. As well as features of the individual, PIMS need to measure and model the world 'out there'. From train times to supply chains, from financial risk to social status, the challenges of capturing and marshalling this information on behalf of the individual are huge. Taken to their logical conclusion, these challenges could be seen as part of the more general projects of artificial intelligence and the semantic web [25].

Some of the important design challenges faced by PIMS in this context have been explored by Mireille Hildebrandt [26], who identifies in particular the asymetry between the epistemic possibilities of data science and the bounded rationality of the human condition. The solution lies in new kinds of 'human-machine interfacing' capable of making transparent the 'back end' of the new computational systems which drive 'marketing, law enforcement and the whole plethora of automated decision-making systems that co-define our lifeworld' ([26] p. 239). The PIMS outlined here are a prime example of such systems, albeit exclusively working on behalf of individuals rather than organisations. In order to provide their users with meaningful agency as an ideal observer would, PIMS will need to integrate the kind of transparency capabilities Hildebrandt describes. Otherwise they risk being black boxes into which individual agency disappears, rather than tools which transparently and comprehensibly extend it. In other words, PIMS need not only to emulate the ideal observer, but also demonstrate how they do so in a way that their users can understand.

[24] See related research on nudging better privacy decisions [6].

[25] The notion of 'nudging oneself' has been previously explored by the author [9], and by the RSA *Social Brain project* which uses the term 'steering' [1].

A related challenge is that of gaining trust. By emulating the ideal observer, and making their workings transparent through new human-machine interfaces, PIMS will go some way to ensure trust.[26] But beyond this, why should individuals trust PIMS any more than traditional businesses? Despite their marketing rhetoric, promises, and manifestos, are PIMS necessarily more likely to serve consumer interests than any other consumer-facing brand? Not necessarily. Broadly speaking, there are two business strategies available to PIMS, which can be characterised as *locust* and *bee* strategies [27]. *Locusts* attempt to extract as much value from a market as possible, taking advantage of any power asymmetries they find themselves on the right side of. Consumers might well fall on the wrong side of such asymmetries, and become subject to the extractive behaviour of PIMS. *Bees*, on the other hand, enrich themselves only as a by-product of enriching others.

Building their business model and organisational structure around *bee* strategies may be the best way for PIMS to gain trust with consumers. In some cases, such strategies might be embedded into the legal structure of the organisation; for example, the personal data store provider Mydex operates as an asset-locked community interest company to ensure that their business purposes remain aligned with the interests of their users [28]. This suggests that PIMS may not just be traditional businesses providing a new type of service, but perhaps an example of a new kind of organisation altogether.

3.5. The Ideal Observer and the Emerging PIMS Market

The previous subsection invoked the rather abstract ideal observer as a requirement for good PIMS design. In this final section, we return to the sample of real-life PIMS introduced in Section 1.1, to assess how closely they approximate this ideal. There are several dimensions along which their success can be measured.

One dimension is the range of values they are capable of being sensitive to. Promoting economic benefits like financial savings is an important (perhaps the most important) consideration for PIMS to be able to factor in, but economic benefits are not the only consideration that a holistic ideal observer would take into account – there are also wider values deriving from general well-being, prudence, spirituality, ethics, or aesthetics. The ideal observer would consider all the relevant dimensions, weighing them up against each other.

For the most part, existing PIMS lack this holistic perspective, and tend to embody *homo economicus* more than they do the holistic, nuanced ideal observer. The personal finance managers, price alert systems and energy switching platforms all have the explicit aim of saving people money, but are insensitive to other kinds of vaues and goals. However, there are a small but growing number of PIMS which help their users achieve non-financial aims. For instance, EthicalConsumer and GoodGuide help their users integrate ethical considerations into their purchasing decisions. In addition, a number of specialist PIMS not mentioned in the introduction focus on physical and mental well-being.[27]

But with different PIMS promoting different – possibly conflicting – aims, individuals may be left struggling to balance one abstract value against another. Few PIMS have been found which combine prudential, ethical and other considerations into one

[26] The application of the concept of the ideal observer to the creation of trustworthy digital systems has been suggested previously by Kieron O'Hara [3].

[27] These include services like QuantID and uMotif, which provide personal health and fitness-tracking dashboards.

decision-support service. Such a service would need to enable individuals to balance one valuable end against another. For instance, a decision-support service which helps consumers select a bank might factor in both the quality of their customer service and their record on ethical investment. Or a travel management service might suggest the least stressful means of transport rather than simply the cheapest or fastest.

A second design requirement for ideal observer-oriented PIMS is that they create a meaningful deliberative route from the individual's current motivations to a more developed and considered set of values. In other words, PIMS should be sensitive to an individual's expressed or revealed preferences, but be able to steer them towards more sophisticated decision and action. While many of the PIMS surveyed above simply take their users' stated preferences at face value, some do make attempts to develop those preferences to become more informed, coherent and sophisticated. For instance, they help their users overcome various cognitive biases, introduce previously unknown sources of information, or help individuals understand the importance of previously ignored factors. For instance, financial managers can help their users exercise good fiscal strategies by overcoming their cognitive bias for immediate gratification or loss aversion. Some of the more sophisticated product-finding services don't just help people select the cheapest or most highly rated product, but help them understand which factors and features are worth paying for in a product and why. Similarly, services like Ethical Consumer help people to choose products based on ethical factors they say they care about, but also provide educational reports which inform their users about what specific things they *should* care about given their stated ethical commitments. There is much room for improvement in these deliberative processes, but they mark the beginnings of a PIMS model which could help individuals arrive at decisions and actions which more closely emulate their ideal observer.

Finally, PIMS need to make their inner workings transparent by way of intuitive, understandable human-machine interfaces as described above. In this dimension, many PIMS fall short. For the most part, the criteria used to drive decision-support services are over-simplified, unexplained, or hidden in confusingly worded terms and conditions. This is particularly problematic in cases where recommendations might be influenced by commercial relationships with providers – for example, where a price comparison website favours results from providers with whom they have commission arrangements. A lack of transparency may go unnoticed in the short-term; most new users will not have the time or inclination to assess such systems in detail, especially where the algorithms used are complex. But in the long run, opaque processes may reduce trust in PIMS. Transparency is a prerequisite for users to be able to verify that PIMS are extending their agency, rather than undermining it.

4. Conclusion

This chapter has explored what it would take for personal information management services to be genuinely empowering for individuals. It has argued that PIMS shouldn't be considered primarily as a new market in themselves, but rather for their potential to change how markets work in general. It has identified the complexity of modern consumer markets as a problem for which PIMS might be a general solution, and introduced some guiding principles for their design. In short, by helping individuals embody their

'ideal observer', PIMS can offer genuine empowerment without sliding into paternalism and manipulation, or undermining autonomy and integrity. Finally, a brief assessment of existing PIMS suggests that while some progress is being made towards embodying these principles, there is still significant room for improvement.

It may seem unusual that in a discussion about personal data and empowerment, the concept privacy has hardly been mentioned. This is an intentional attempt to carve out a distinct conceptual space for the notion of personal data empowerment as distinct from privacy. It is true that power and empowerment do frequently enter into existing discourse about privacy; primarily in terms of the power of individuals to control access to their personal data [29], and the intersection of privacy with existing power dynamics [30]. These discussions are important and worthwhile, but the notion of privacy they employ is distinct from that of personal data empowerment. Personal data empowerment refers to a situation where individuals are able to use data, including their own personal data, for their own purposes, to achieve personal aims and goals. As I have argued, it is the answer to a broad challenge; namely, restoring and enhancing the capacity for meaningful agency in an increasingly complex world.

Whilst the rather lofty goal of embodying an abstract philosophical ideal is unlikely to feature in entrepreneurs' business plans any time soon, it nevertheless offers a framework for analysing market-driven empowerment through PIMS. The ideal observer is just that – an ideal – which it is impossible to embody fully, but which it is possible to strive towards. Personal information management services have the potential to bring individuals far closer to this ideal than they could ever have come on their own.

Acknowledgements

In addition to conducting independent research in this area, the author has been employed by industry analysts Ctrl-Shift as a researcher, and contributed to a number of their reports, which are referenced within. Although it doesn't necessarily reflect the opinion of Ctrl-Shift or its clients, this chapter has benefited significantly from these reports, and builds on their concepts and findings. The author would therefore like to acknowledge and thank the rest of the Ctrl-Shift team for this work, in particular Alan Mitchell for his insight and input in an initial draft. The final draft also owes much to the helpful feedback from the editors and reviewers.

References

[1] M. Grist, "Steer: Mastering our behaviour through instinct, environment and reason," RSA, 2010.
[2] J. Mokyr, "The European Enlightenment, the Industrial Revolution, and Modern Economic growth," vol. 2007, no. June. Max Weber Lecture, European University, 2007.
[3] K. O'Hara, "Trust from the Enlightenment to the Digital Enlightenment," *Digit. Enlight. Yearb. 2012*, pp. 169-183, 2012.
[4] D. Kahneman, *Thinking, Fast and Slow*. Macmillan, 2011.
[5] W.J. Baumol, *The free-market innovation machine: Analyzing the growth miracle of capitalism*, 2002.
[6] R. Balebako, P.G. Leon, H. Almuhimedi, P.G. Kelley, A. Acquisti, L.F. Cranor and N. Sadeh, "Nudging Users Towards Privacy on Mobile Devices."
[7] Ctrl-Shift, "Next Generation Intermediaries," January, 2014. Available: https://www.ctrl-shift.co.uk/research/product/85

[8] N. Ashraf, C.F. Camerer and G. Loewenstein, "Adam Smith, Behavioral Economist," *J. Econ. Perspect.* **19**(3) (2005), 131–145.
[9] R. Binns, "Nudge Yourself," 2013. [Online]. Available: http://www.reubenbinns.com/blog/nudge-yourself/. [Accessed: 02-Jul-2014].
[10] Ctrl-Shift, "Personal Information Management Services?: An analysis of an emerging market," no. June, 2014. Available: https://www.ctrl-shift.co.uk/research/product/90
[11] D. Chadwick, "An Advanced Policy Based Authorisation Infrastructure," 2009.
[12] P.M. Schwartz, "Property, Privacy, and Personal Data," *Harv. Law Rev.* **117**(7) (2004), 2056.
[13] M. Hildebrandt, K. O'Hara and M. Waidner, "Digital Enlightenment Yearbook 2013: Introduction," *Digit. Enlight. Yearb. 2013 Value Pers. Data*, pp. 1–25, 2013.
[14] H. Grassegger, "Das Kapital, Bin Ich," vol. Kein & Abe, 2014.
[15] K. Crawford and D. Boyd, "Critical questions for big data," *Information, Commun. Soc.*, no. October 2013, pp. 37–41.
[16] A. Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations*, 1776.
[17] G.J. Stigler, "The successes and failures of Professor Smith.," *J. Polit. Econ.*, 1199–1213.
[18] H. Simon, "Theories of bounded rationality," *Decis. Organ.* (1972), 161–176.
[19] A. Smith, *The Theory of Moral Sentiments*, 2010th ed. 1759.
[20] R. Firth, "Ethical absolutism and the ideal observer," *Philos. Phenomenol. Res.* **12**(3) (1952), 317–345.
[21] B. Williams and A. Quinton, *Utilitarianism, For and Against*, 1973.
[22] P. Railton, "Moral realism," *Philos. Rev.* **95**(2) (1986), 163–207.
[23] C. Sunstein and R. Thaler, *Nudge: Improving decisions about health, wealth and happiness*. Yale University Press, 2008.
[24] C. Sunstein and R. Thaler, "Libertarian paternalism is not an oxymoron," *Univ. Chicago Law Rev.*, 2003.
[25] B.T. Berners-lee and J. Hendler, "The Semantic Web," *Sci. Am.* **21** (2002).
[26] M. Hildebrandt, "Profile transparency by design: Re-enabling double contingency," in *Privacy, Due Process and the Computational Turn*, M. Hildebrandt and K. De Vries, Eds. Routledge, 2013.
[27] G. Mulgan, *The locust and the bee: predators and creators in capitalism's future*. Princeton Univ Press, 2013.
[28] W. Heath, D. Alexander and P. Booth, "Digital enlightenment, Mydex, and restoring control over personal data to the individual," in *Digital Enlightenment Yearbook 2013: The Value of Personal Data*, M. Hildebrandt, K. O'Hara, and M. Waidner, Eds. IOS Press, 2013, pp. 253–269.
[29] A.F. Westin, *Privacy and Freedom*, vol. 97, no. 4. Atheneum, 1967, p. xvi, 487 p.
[30] J.E. Cohen, "What Privacy is for," *Harv. Law Rev.*, vol. 14, no. forthcoming, pp. 1–24, 2013.