CHAPTER

1 Introduction

Is data the new oil? Data has taken on a critical role with the rise of the digital economy. In the past decade, companies with data at the core of their business models have come to dominate the rankings of the world's most valuable corporations. The literature that studies the economics of data spans separately growth, privacy, competition, inclusion, and financial stability aspects. This paper provides an analytical review that integrates these perspectives and assesses the implications of data for macroeconomic growth, equity, and stability. Its contribution is to describe the main trade-offs facing policymakers as they design data policy frameworks for the increasingly complex global data economy.

The proliferation of data in the economy presents a tremendous opportunity to boost growth through efficiency and innovation. But to do so without compromising other objectives, we argue that current policy frameworks must be modernized to tackle four growing challenges. First, data markets are opaque and may be leading to too much data collection and too little privacy. Rights and obligations over data must be clarified for the market to function efficiently, and the way in which these are assigned will impact growth and equity. Second, incumbents have an incentive to hoard data, potentially stifling competition and reducing the social benefits that could flow from wider access. A range of policies can be deployed to encourage data sharing that can promote competition and innovation. Third, it is unclear that companies are doing enough to protect the data they hold, creating risks to stability that should be mitigated with measures to ensure that all market participants invest adequately in cybersecurity. Finally, without some coordination across countries, there is a risk that global data markets could become fragmented, impeding potentially large gains from cross-border data activity including trade and finance.

This page intentionally left blank

CHAPTER

2 Context and Framing

Data has long been of value in economic activity. For a colonial bank to extend a loan to a farmer, it would consult the data in the local registry to learn how much land its potential borrower owned. Ancient states conducted censuses to gather information on their subjects to facilitate the task of levying taxes. Many market activities have been recorded to build trust and provide assurance on the exchange of goods and services. The collection of personal data has always involved a trade-off between respecting the individual's desire for privacy—including from government—and reaping the commercial and social benefits that can be derived from its collection and dissemination.

What is new about data that requires policymakers to rethink its economic implications? Two recent technological trends are widely recognized as having led to an explosion in the economic relevance of data. First, technological progress has drastically reduced the costs of collecting and storing data. Widespread digitalization leads to more data being produced as a byproduct of economic and social activities, including aspects of human interactions and experiences that used to be conceived as being entirely qualitative. Second, advances in analytic techniques have allowed for more advanced processing to extract greater value from available data. General purpose technologies including artificial intelligence and machine learning have pushed the use of data across sectors, with prediction algorithms deployed to develop self-driving cars, identify promising new drugs, deliver targeted advertising, and to improve the efficiency of operations. For several of the world's most valuable publicly traded firms, data is central to their highly profitable business models.[1]

[1]In their report to the US Securities and Exchange Commission from July 2019, Alphabet (GOOGL) reported that advertising revenues—generated by the company's data-driven ad targeting services—reached $32.6 billion in the latest quarter, making up 83.7 percent of total revenue.

These trends are changing how consumers, companies, and policymakers measure and analyze the economy. Previous IMF work has studied the implications of big data and digitalization for the compilation of economic statistics (IMF 2018a) and for real-time policy making (IMF 2017b).[2]

This paper focuses on the economic characteristics of data and their implications for macroeconomic growth, efficiency, and stability. Although the proliferation of data is a recent development, the economic literature contains many relevant insights that we draw on. At the outset, we should emphasize that while the literature offers a rich set of qualitative mechanisms and trade-offs, their quantification remains at early stages, as are efforts to describe optimal policies by incorporating multiple mechanisms into unified models.

There is a long tradition of study in economics on the importance of imperfect information—especially when it is held asymmetrically—as a key friction undermining allocative efficiency and compounded by search and transactions costs (Stiglitz 2002). Likewise, much has been written about the role of knowledge and ideas in driving economic growth (Arrow 1962; Romer 1990). An emerging literature has focused directly on the economics of data and digitalization more broadly. Jones and Tonetti (2018) and Farboodi and Veldkamp (2019) present growth models with data in the production function and study its implications for long-term growth. Acquisti, Taylor, and Wagman (2016) provide a synthesis of the insights uncovered in the literature on the economics of privacy, which has a rich tradition going back to Posner (1981). And more broadly, Goldfarb and Tucker (2019) provide a review of how digital technologies are changing economic decisions by shifting the costs of search, duplication, transportation, tracking, and verification.

With the literature on data large and growing rapidly, the main contribution of this paper is to integrate perspectives to inform the design of data policy frameworks. We focus on the macro-financial implications of data proliferation through its impact on efficiency, equity, and stability. We start by asking: What does data do in the economy? We consider two functions that have been emphasized in the growing literature on the topic. First, data is an *input* in the modern production function that firms combine with factors such as labor, capital, land, and oil to produce a wide range of goods and services, and to innovate. Second, data *shifts information* across economic agents, with implications for efficiency, equity, and competition.

We then argue that data has three economic characteristics that create important challenges for public policy.

[2]The IMF (2017a) has also explored how the combination of big data with artificial intelligence, distributed computing, and cryptography could change the provision of financial services, with a special emphasis on cross-border payments.

First, data is *nonrival*: While using oil means others can no longer use it, the same data can be used by many. Like a new idea, society will benefit most from data when it is widely shared, because more users will be able to use it to increase efficiency and innovate. But while technology makes data nonrivalry possible, policies and private decisions affect whether it will be so in practice. Under current policies, it is unlikely that private firms have incentives to grant competitors access to the data they have collected, such that data hoarding practices may be limiting market contestability and the social benefits that could be derived from data.

Second, data involves *externalities*: The collection, sharing, and processing of personal data by one agent imposes costs on others by affecting their privacy. An implication is that a market for data lacking sufficient user control rights—where data collectors do as they please with the data they collect—is likely to lead to excessive data collection and too little privacy. A related policy challenge is to clarify the rights and obligations of participants in data markets.

Third, data is only *partially excludable*: The storage of data on interconnected systems means that controlling access to data requires continuous investment to prevent its loss through cyber-attacks. A key policy question is to what extent private data collectors and processors have adequate incentives to invest in protecting their data, particularly in the case of individual data about others. There is an emerging consensus that private reputational effects are insufficient, and policy measures are needed to ensure that sensitive data is protected adequately.

Effective data policy requires an integrated perspective to balance competing objectives: promoting growth and competition through data access, ensuring incentives exist for data to be collected and processed, promoting stability by adequate investment in cybersecurity, and ensuring that individual privacy preferences are respected. In most cases, this will require cooperation at the national level across agencies that may not have traditions of interaction: consumer protection and privacy agencies, competition authorities, ministries of finance and economics, statistics offices, central banks, and financial regulators.

It is natural that national priorities will vary across objectives, such that there can be no one-size-fits-all data policy framework. However, there is also a risk that the global data economy becomes fragmented. Because data is inherently mobile, such a scenario could stifle the potential benefits of international trade in data. We argue that international coordination is needed to achieve the minimum data policy principles that are compatible with productive cross-border data economies.

This page intentionally left blank

CHAPTER

3 The Building Blocks of a Market for Data

This section lays out the basic ingredients needed for thinking about the economic decisions involved in a stylized market for data. When we refer to data, we refer to the factual representation of a characteristic, action, or natural occurrence. Data can be quantitative or qualitative in nature, and may be stored on analog (that is, paper, stone tablets) or digital media.

Data is a form of information, and a rich literature on the economics of information—and particularly on incomplete or imperfect information—thus offers many useful insights for thinking about the economics of data (Stiglitz 2002 provides an overview). We follow Jones and Tonetti (2018) in distinguishing data from ideas, which are another form of information. The intuition is that an idea is information that provides a set of instructions for completing a task, such as a recipe for producing a drug or a blueprint for building a machine. In contrast, data is an ingredient that can be used in the completion of a task, more akin to a raw material.

This remains a broad definition of data and necessarily makes much of the discussion in the following sections relatively abstract. Data is heterogeneous along many dimensions, including its subject(s), timing, format, and quality, and some of these features will affect its economic characteristics and implications. A particularly important dimension is whether the data can be mapped to an individual person (personal data) or to a firm's operations and trade secrets (commercial data). Another important dimension is the amount of analysis data embeds. We understand "data" to be recorded as observed or to have undergone a degree of processing to standardize it—examples could include recorded facts about a transaction, footage from a surveillance camera, locational information, or a person's online browsing history. In contrast, digital content such as a software program, a file containing a motion picture, or a synthetic credit score all embed substantial effort and proprietary knowledge to produce, and are not the main focus of the analysis in this paper.

Supply: The Decision to Produce Data

The economic agent who collects data and stores it is referred to as the *data collector*. In doing so, the data collector incurs costs. Some of these costs are fixed, such as the installation of sensors or technological infrastructure (consider a passenger transport company that installs a system of smart payment cards and scanners), whereas others vary as a function of the quantity of data collected, such as storage capacity, or the payment to any labor that was involved in creating the data (consider a clinical trial where participants are compensated for subjecting themselves to an experimental treatment). Farboodi and Veldkamp (2019) emphasize that marginal costs of data collection are very low where data is generated as a byproduct of economic activity, such as production or transaction data (an airline that tracks the number of passengers on its flights and how much each paid). But even in these cases, the storage and protection of data still represents a variable cost for the data collector that can be substantial.

In the case of personal data—which can be mapped to an individual—the person whose characteristic has been recorded is referred to as the *data subject*. When collecting such data requires the active participation of the data subject, the data collector may incur a cost to compensate the data subject. For instance, a supermarket may offer customers discounts for using a loyalty card that tracks their buying habits. Many digital platforms operate two-sided businesses, where data is collected as part of a barter transaction in exchange for access to a virtual service and is then used in a transaction with a separate client, such as an advertiser. Although no money changes hands between the data subject and data collector in these cases—it is a barter exchange of data for a service—providing the virtual service represents a cost for the latter.

The cost of obtaining the data subject's consent to collect their personal data is closely tied to their preference for privacy, which reflects the disutility caused by having their personal information shared with others.[1] A rich literature on the economics of privacy has studied the theory and empirics of individual preferences for privacy (Acquisti, Taylor, and Wagman 2016 provide a comprehensive survey). These preferences naturally vary depending on the individual trait or action being represented in the data—supermarket shopping habits may be perceived as relatively innocuous to share with others, in contrast to health files or criminal records—and will also vary across individuals and cultures.[2] The cost will also depend on the data subject's awareness of the transaction and the saliency with which they perceive the privacy implications of their decision, which may materialize in the future.

[1]As we will discuss in the following section, a failure to recognize the right to privacy can be conceived as an economic externality.

[2]Winegar and Sunstein (2019) present survey evidence that individuals demand significantly more money in exchange for granting access to their health data than to their demographic data.

Economies of scale can take place when the average cost of collecting data decreases as more is collected. This will tend to be the case when initial fixed costs are large but marginal costs are very low, and this is common in firms that invest in extensive digital infrastructure to collect data. A similar dynamic is observed in platforms—including social networks and payment systems—where demand-side network externalities are such that participation in the network becomes more attractive to marginal users as the number of users (and the data they generate) grows. This means that scale increases the marginal user's willingness to surrender their data in exchange for access, thus lowering the platform's barter cost of acquiring more data (Furman 2019).

The cost of data collection is naturally a function of technological developments, and recent changes have altered the cost curve in two ways. First, the rapidly falling cost of digital sensors (including cameras, microphones, global positioning systems, and accelerometers), storage technology, and the proliferation of digital economic activity have drastically reduced the cost of collecting and storing data.[3] Second, the cost curve faced by data-collecting firms has been shifted down by the development of third-party data intermediaries such as cloud storage services, which transform the fixed costs into variable costs. These intermediaries thus tend to reduce the economies of scale faced by data collectors, and do so by achieving very large scale themselves.

These changes to the costs of collecting and storing data have contributed to an increasing amount of data being collected about a growing number of phenomena, by a growing number of data collectors.

Demand: The Economic Value of Data

If creating and storing data is costly, why do so? We now turn to the demand for data, which is determined by how data is used in the economy. A data collector will only decide to record data when the costs of doing so are smaller than the revenues they expect to derive from it. In the early phase of digitalization, many firms simply discarded most of the data generated by their operations and users' activity, assuming it was not of significant value.

The literature has focused on two roles of data in the modern economy.[4] First, data is an input into the production of goods and services, and contributes to innovation and efficiency. Second, data creates and shifts

[3]In many cases, data is generated as a byproduct of other economic activities. Even then, its storage, management, and ensuring its security all represent costs to the collector.

[4]Data may also be collected for use as a consumption good itself, as in the case where users use smart devices to monitor their biometric activity—the number of steps they have taken in the day or the length of their sleep cycles—often as a hobby and sometimes with medical ends in mind. We focus here on data collected for commercial purposes.

information across economic agents, affecting strategic interactions and information frictions.

Data as a Factor of Production

The first function of data is as an input in the production of a good or service. In this function, data enables the creation of knowledge, which can be directed toward the ongoing production of an existing good (for instance, observational data needed to run a weather forecasting service or provide a targeted ad) or in the development of a new product or service. Deriving value from data as an input requires costly processing and analysis, which is provided by complementary skilled labor. The agent that aggregates and analyzes the data is referred to as the *data processor*. The data collector and processor may in practice be the same company or individual, but this need not be the case when data can be traded.[5]

Data analysis can also be used in innovation, as new insights extracted from data can lead to the development of new products or services. More recently, the proliferation of big data and the development of machine learning algorithms have enabled data analysis to address increasingly complex problems. Data has thus come to represent a necessary input into the development and production of a wide range of new products. For instance, cars equipped with sensors may record the actions of drivers as they navigate city streets, building up a massive data set of human decisions in the face of various situations. Patterns in this data can then be analyzed using machine learning algorithms to predict and mimic human decision-making in complex road environments, which may then enable the production of a safe self-driving car.

Data Creates Information and Shifts It Across Agents

The second function of data in the economy is in creating and shifting information. Varian (2018) emphasizes that the firm's collection of data about its own operations facilitates a process of learning by doing. For instance, analysis of data can reveal insights that firms use to modify and improve their business practices. The longer the firm operates, the more corporate data it generates as a byproduct of its operations, and the more it can learn from its own past production decisions, allowing it to become more efficient.[6] This is the approach taken by Farboodi and Veldkamp (2019), who model data as

[5]Note that under the European Union's General Data Protection Regulation, the act of data collection is also considered "processing."

[6]This function is not restricted to the private sector. Public institutions also analyze data to better understand the impact of their policies.

providing information that reduces uncertainty about random variables that are relevant for production.

Personal and corporate data embeds information about economic agents—including consumers or firms—so access to it shifts information asymmetries in markets these agents participate in. When access to data serves to reduce information asymmetries between buyers and sellers, it can potentially lead to more efficient economic transactions. For instance, a seller with access to data about the characteristics of potential consumers—such as their interests and buying habits—can deliver a more personalized good or service, such as an advertisement for a product that they are more likely to find desirable. Likewise, consumers with data on characteristics of potential products can make more informed buying decisions by more accurately assessing how products fit their needs. This can include data about product reliability (for instance, data on reviews from past customers), popularity (for instance, the use of services, accessories, and clothing by influencers on social media), and rapid comparisons with how they compare with competing products (for instance, through retail aggregation services).

The data collector's willingness to pay for data depends on many factors, including how much existing information asymmetry they face with their customers and competitors, their degree of market power, and the size of the market. Assessing how valuable personal data can be turns out to be quite difficult, even for the agents that have direct incentives to do so. For instance, advertisers spend large sums on individual or pooled data about online users, on the premise that displaying a more targeted ad will make it more likely to generate sales. However, an incipient literature quantifying the effectiveness of this practice suggests that, although the gains from targeting ads do appear to be statistically significant, their causal impact on sales appears to be economically modest and inferior to the outlays spent on targeting (Marotta, Abhishek, and Acquisti 2019).

Whereas more information can increase economic efficiency, acquiring information that others do not have provides a strategic advantage, potentially making some groups worse off. Acquiring data may thus generate considerable commercial (private) value for a data collector, but without necessarily increasing social welfare. If a firm enjoys market power, then gaining data about their customers' personal characteristics—say, their income or wealth—can allow them to implement price discrimination strategies that extract the consumer's surplus. For instance, an airline in a poorer country may offer higher-priced fares to foreign consumers, who are likely to be richer and willing to pay (what microeconomists refer to as third-degree, or group, price discrimination). Likewise, contractors may adjust the price they charge for services if they know the value of the customer's home. In the extreme, a

monopolist seller with precise information on a customer's income and tastes might be able to implement first-degree price discrimination, thus extracting all surplus from consumers.[7]

Data's function in shifting information across agents that participate in financial markets has been the subject of a well-established literature. The provision of financial services—from savings to intermediation and payments—has always relied upon data to record transactions and reduce information frictions. As noted in "The Bali Fintech Agenda" (IMF 2018b), access to customer data is an important input into the business of financial intermediation and a key part of the promise that technology offers to the provision of financial services.

An important function of financial intermediaries is to channel idle savings to productive investment projects and consumption opportunities. To do so, lenders require data on potential borrowers to gauge their creditworthiness and to monitor their performance after a loan has been extended. Incomplete information between lenders and borrowers prevents the efficient allocation of credit because of adverse selection, which represents a particularly relevant friction in developing economies.[8] Stiglitz and Weiss (1981) show that, when lenders do not have full information about individual borrowers' ability and willingness to repay a loan, they are likely to ration credit—that is, some borrowers will not be offered a loan at any interest rate and will thus be excluded from the financial market. Access to more granular user data can reduce information asymmetries, holding the promise of reducing lending costs and expanding the availability of credit.

The Price of Data

Data is not homogeneous but rather is differentiable based on a large number of attributes: who it describes, when and where it was collected, how structured is the data, and if it can be merged with other varieties. Within narrow classes of data varieties, it may be possible to define meaningful markets for data with a single price, which will vary significantly across varieties and over time.

The market for data involves agents trading data at a price that would reflect some of the trade-offs we have described. But what price should we expect

[7]Ezrachi and Stucke (2016) argue that data-based "behavioral discrimination" has the potential to complement perfect price discrimination strategies by using personalized emotional cues to influence consumer preferences.

[8]Adverse selection refers to the fact that the borrower's willingness to accept a very high interest rate signals to the lender that they are unlikely to repay the loan. The higher the interest rate offered by the lender, the riskier the pool of borrowers willing to accept it.

to clear the market? The degree of complementarity or substitutability across data varieties will have an incidence on their price. Different varieties of data may be complementary when merged, for instance, if doing so enables the identification of relationships that could not be inferred from analyzing the data in isolation. Other data types may be substitutes for each other, particularly if they share common attributes (for instance, they describe similar populations along a common dimension of interest). For instance, driving data produced by one driver may be equally informative for the self-driving car problem than data produced by another driver.

Economic Characteristics of Data

Recent developments in information technology have made three economic characteristics of data more salient. Each, in turn, has implications that can undermine economic efficiency, equity, and stability in the absence of appropriate policies. We describe these characteristics briefly and will turn to their implications and policy responses in the following sections.

Nonrivalry

The digitalization of data and the ability to transfer it across networks has made data increasingly nonrival by virtually eliminating the costs of duplication and transfer. In contrast to most other goods—but like other types of information—one agent's use of data does not diminish the ability of others to use it, even simultaneously. This powerful characteristic holds the promise of enabling increasing returns to scale and scope through data sharing, as we will describe in the following section. Where data proliferates in production, bigger firms will tend to be more productive. Whether these returns to scale are exploited within firms or across firms will have a bearing on market structure and competition.

Varian (2018) emphasizes that data's nonrivalry means that it is rarely bought and sold like other goods or services, but rather is licensed for specific uses. The relevant economic question of control over data is then more of *access* than of *ownership*. An implication is that, before we can have an efficient and responsible market for data, we need to agree on who controls it—who will have access to it, and who won't—so that its benefits can be derived and shared fairly and traded off in a considered manner that appropriately balances costs, security, and privacy.

Jones and Tonetti (2018) propose a growth model with data as an input in the firm's production function, which is created as a byproduct of economic activity. But unlike other inputs—labor, capital, land, or oil—they emphasize

that data is nonrival. An important implication of the nonrivalry of data is that, from a social perspective, it is desirable for data to be widely shared. Jones and Tonetti (2018) warn that this outcome will generally not be compatible with the private incentives of data processors, however, who tend to hoard data they collect to gain an advantage over their competitors and avoid competition from new entrants. A market outcome may thus give rise to suboptimal concentration and less contestability, as incumbents protect their data advantage over potential competitors by limiting access to it. This will substantially lower the growth benefits from data with respect to what could be achieved under wider sharing.

Privacy Externalities

The ability to share data over networks and to make it public to a global audience has increased the saliency of its privacy implications. Important benefits can accrue to the data subject from sharing their personal data, including the provision of innovative services and more customized products. But when the data subject is unaware, decisions made about their personal data can give rise to an externality: private decisions about whether to collect, process, or share personal data have a bearing on the economic well-being of the data subject, who may not be compensated. These externalities are often negative, with the use of individual data imposing disutility on data subjects through two channels. A direct effect is that the sharing of data inherently undermines the data subject's preference for keeping their personal characteristics or actions private. In addition, individual data may be used strategically by economic agents that acquire it to extract rents from the data subject.

For the market to function efficiently, the benefits of revealing data must be weighed against the harm that can come from reducing privacy. As Acquisti, Taylor, and Wagman (2016) underscore, privacy should not be understood narrowly as preventing the sharing of personal information, but rather as giving the data subject control over what they share. To the extent that privacy is not internalized in the economic decisions of data collectors and processors, the market will tend toward the collection of excessive personal data and insufficient protection of privacy. For the market for data to internalize this externality, the rights of data subjects must be adequately attributed.

A key question is whether, given the substantial benefits to consumers and markets from revealing personal data, granting strict user control rights would lead them to stop sharing their data in most cases, which may make some services unviable and stifle future innovation. In regimes that grant users control over their data, it is important for data processors to provide consumers clarity on the value of services being offered in exchange for granting data access. Where the value of the service is perceived to be small, data

subjects may demand financial compensation to engage in the transaction. Alternatively, a few jurisdictions have considered instituting some form of data dividend scheme whereby a portion of rents garnered by data processors are shared with users automatically.

Can technologies such as anonymization provide a win-win by enabling the benefits of data access while maintaining adequate privacy? In several applications, data analytics can provide valuable insights without data being individually identifiable. Consider training artificial intelligence to drive an automated car, to recognize images of specific objects, to give a medical diagnosis based on an x-ray or blood test, or to study the effects of new pharmaceuticals based on anonymized data. All these applications rely on huge amounts of data to train the algorithms, but do not require that the data be linked to an individual. But in many applications, the value of data is substantially reduced through anonymization. One reason is that although anonymized data can still be used as an input to certain analytic tasks, it generally no longer reduces information asymmetries. In many instances, keeping data private will involve an efficiency cost. Posner (1981) emphasizes the efficiency costs of excessive privacy protection, because individuals that withhold material information can inflict substantial harm on their counterparts.

How large is the data privacy externality? On the one hand, reputation effects may create private economic incentives for a firm to ensure their services are designed to respect and protect the privacy of its users. For a data-based company, the risk of being perceived as lax in their protection of user data—either from a data breach or by the revelation of misleading privacy policies—may lead users to limit sharing their data with the firm.

On the other hand, valuing privacy is inherently difficult, even for one's self. Research in the literature on privacy has identified an apparent *privacy paradox*, whereby individuals place a much lower value on their privacy in their actions than they do when asked to place a subjective value in surveys. A common example are the electronic disclosures ("I agree") that online platforms require their users to accept prior to using their services. Whereas a large percentage of people tell surveyors they are very concerned by a company sharing their private information, almost all willingly grant their consent to do so in exchange for the most basic of "free" online services.[9] But an important challenge here is the extent to which consent given is truly informed and the data subject has agency over the choice they are expressing, with instances cited of manipulation by data processors. In discussing survey

[9]Some scholars argue that this is not a paradox at all, but rather reflects the difficulty of evaluating these types of trade-offs, which are inherently intertemporal and opaque. For instance, a user must decide if they are willing to sacrifice the tangible benefit of using a webpage today versus the potential cost of having their privacy compromised in unspecified ways at an unspecified time in the future.

evidence on consumers' stated valuations of privacy, Winegar and Sunstein (2019) argue that information deficits and behavioral biases make these valuations uninformative about the true economic value of privacy.

Partial Excludability

So far, we have assumed the question of access and control of data is a choice. However, technology has made data inexpensive to duplicate and transfer through interconnected systems that are vulnerable to cyber-attacks. In contrast to the stone tablets of antiquity, the digitized data used in the modern economy is only partially excludable. We may stipulate that a data transaction is to involve only two parties, but technology may make this difficult and costly to enforce in practice. Agents that make economic decisions about data must consider the risk that it will become available beyond the intended counterparts.

To be sure, data can be made almost perfectly excludable by storing it offline on a secured, isolated system. But this will drastically reduce the commercial and social value of the data, particularly in applications where real-time data is needed, such as in logistics, financial transactions, and targeted online advertising. In practice, there may be a continuum of options for managing excludability of data reflecting decisions made by economic agents as well as technological factors. They may incur a variable cost to protect data by hiring engineers to secure their networks, or by training their staff to use good cyber hygiene. Further costs can be incurred by implementing encryption or by partially anonymizing the data. Data collectors and processors thus face an important economic decision on how much to invest in cybersecurity to protect their commercial data and the personal data they hold about consumers.

Private data collectors and processors will have incentives to incur costs to secure their data from corruption or misuse. Restricting access to a strategic asset is important to maintain a comparative advantage over competitors. Avoiding harm to their reputation as good stewards of their customers' data may be important to support demand for their services. Finally, the harm caused to their clients by the misuse of their individual data could create a liability for the firm.

However, the private incentives to invest in data security are unlikely to lead to socially optimal levels of investment in cybersecurity (Kashyap and Wetherilt 2019). Nonrivalry and partial excludability make private contracts difficult to enforce, because the harm caused by misuse of data is difficult to trace to a specific breach. With incomplete contract enforceability ex post, an efficient digital economy requires agents to trust that their information will be adequately protected by counterparties (Organisation for Economic

Co-operation and Development 2015). Perceptions of inadequate privacy or insufficient cybersecurity thus involve an externality, because the investment decisions of individual agents will affect overall trust in the economy's data security. By reducing trust and thus the willingness of users to share their data, one data set being mishandled may cause more harm to the system than the sum of the direct harm caused to each of the data subjects.

This page intentionally left blank

CHAPTER

4 Macroeconomic Implications of Data Proliferation

We have described the broad features of the market for data and some of the economic characteristics of data that can introduce frictions into these markets. We now turn to the implications of the proliferation of data, which arises from a technologically driven fall in the cost of collecting data and from improvements to the analytical tools needed to analyze data and extract valuable knowledge from it. We discuss how these developments are expected to impact macroeconomic outcomes such as growth, equity, and stability.

Growth

As we have described in the previous section, data has the potential to contribute to economic growth through its use as an input in the production of goods and services, and by facilitating firm productivity through learning by doing. By generating information and reducing asymmetries, it can also potentially improve the efficiency of markets, reduce funding costs for borrowers in financial markets, and improve the matching of products to consumers.

Let us begin by considering the case for data proliferation delivering growth by improving the efficiency of the financial markets. Begenau, Farboodi, and Veldkamp (2018) model the proliferation of data as a means of reducing information asymmetries between a lender and borrower, and thus the cost of financing. The more public data is generated by a firm during its operations, the less uncertainty a potential creditor faces about that firm's future cashflow and creditworthiness, and the better the lender can monitor the firm after the loan has been extended. The more data the lender has access to about the borrower, the lower the interest rate they can charge for the loan and the smaller is the share of borrowers who will be denied a loan. This logic is a key component of the business of new financial technology service providers such

as Ant Financial who leverage big data to generate improved real-time credit scoring to provide funding to high-return small and medium-sized enterprises previously rationed out of credit markets because of information asymmetries (see, for example, Hau and others 2019).

Data sharing among financial service providers can attenuate adverse selection and moral hazard problems, and thus increase the provision of credit by reducing information asymmetries. Since the end of the 1800s, credit bureaus have operated as data brokers with whom lenders share information about borrowers on a reciprocal basis—that is, a bank that provides information about its borrowers is entitled to receive any available information about a prospective client.[1] Where private credit bureaus have not emerged spontaneously, many jurisdictions operate public credit registers that mandate the sharing of default information about borrowers. Does sharing data lead to more credit on better terms? Jappelli and Pagano (2002) and Djankov, McLiesh, and Schleifer (2007) present empirical evidence that the operations of data-sharing institutions are associated with deeper and broader credit markets and lower frequency of defaults.

Crucial to data's broader implications for long-term growth is the question of whether accumulating data generates returns to scale and scope. If the value extracted from data were to decrease in its quantity, then data may provide a boost to the level of output, but would not be expected to affect the rate of growth in the longer term. But if instead data were associated with increasing returns to scale, its proliferation could give rise to a source of higher sustained growth.

This remains unsettled and context-specific, with a growing body of work addressing the question. Basic statistical theory states that the power of a statistic tends to grow at a decreasing rate with the sample size. This is the basis for Varian's (2018) observation—citing results from an artificial intelligence competition—that the marginal improvement in precision of machine learning applications decreases in the amount of data that is analyzed.[2] The intuition is that the first million observations are more informative than the next million, and so on.

[1]Pagano and Jappelli (1993) present a model where data sharing about customers can arise endogenously among banks. They argue that the incentives to participate in a bureau are greatest when the bank is faced by many customers on which it has little prior information. However, information sharing may not be in the interest of an incumbent that enjoys some degree of market power, as doing so would subject them to increased competition and face the loss of monopoly profits.

[2]The Stanford dog breed classification project ran a public competition where teams were provided with a data set of tagged photos containing dogs of different breeds. Competitors then wrote machine learning algorithms that could predict the breed from a new image containing a dog.

One argument that can give rise to increasing returns is that firms using data for learning-by-doing can enter a virtuous "data feedback loop." Farboodi and others (2019) present a model in which corporate success leads to the acquisition of more data from its users, which in turn is used to improve productivity and gain market share, which further expands the amount of available data. But increasing returns from such a mechanism may be limited to the firm level, because the number of users generating data cannot be expanded endogenously at the economy level. Whereas it could be a factor shifting market structure toward larger firm size, its ability to generate sustained economic growth would be mitigated by the fact that new entrants may be rendered unable to compete with large data-rich incumbents (Newman 2014). Is there empirical evidence for the data feedback loop? Bajari and others (2019) use data from the Amazon marketplace to test this theory and find that a seller's precision in forecasting demand for their goods grows as a decreasing function of the data at their disposal, implying an upper bound on the gains from scale in data.

Agrawal, Gans, and Goldfarb (2018) argue that the use of data in certain artificial intelligence applications can exhibit increasing returns to scale, particularly in more complex prediction problems. The intuition is that, although the precision of artificial intelligence's predictions may indeed improve at a decreasing rate as it is provided with more data, the usefulness of the algorithm may increase nonlinearly. In many cases, predictions must be extremely accurate to be useful in a commercial application, with threshold effects suddenly turning a useless tool into a profitable product. For instance, the artificial intelligence that generates the predictions needed to operate a driverless car may learn the basics with a modest amount of driving data, but still produces an accident rate that is entirely unacceptable to consumers and regulators. The marginal value of the additional data needed to improve the artificial intelligence's precision past the threshold of acceptable risk—say, the average accident rate of a human driver—would be large and discrete, because it would allow the technology to be deployed in a consumer market.

Goldfarb and Trefler (2018) argue that economies of scale from data used in artificial intelligence can also result from direct network externalities, wherein more customers generate more data, which improves the quality of the product, attracting still more customers, and so on. They also argue that data acquired for a particular purpose can be of value in other contexts, granting data collectors economies of scope in the development of new products. Again, this gives the incumbent firm an advantage over competing new entrants that would have to gather the data or pay the incumbent for access.

Data may affect market structure toward greater concentration by creating barriers to entry that stifle competition. The increasing returns to scale and

scope described so far accrue at the firm level, such that holding a large war chest of data presents a barrier to entry for competitors. This can potentially lead to winner-takes-all dynamics where market concentration rises sharply. Of course, more concentrated market structures are not always negative for social welfare, because they may be consistent with innovation and competitive pricing as long as the market remains contestable in a dynamic sense.

It is conceivable that the strength of network effects—an important barrier to entry created by data—may not last for long. Consider, for instance, how social networks have risen and fallen in popularity as their user bases age. A young user may avoid joining a network specifically because the existing user base includes older relatives and contacts, such that network externalities become negative across cohorts. Observing concentration is thus not necessarily associated with a lack of innovation and growth, as long as dynamic contestability is maintained.

Jones and Tonetti (2018) show that even if we assume that data accumulation itself exhibits decreasing returns to scale—consistent with the arguments and evidence presented previously—it can still give rise to increasing returns in the production function, and thus deliver long-term economic growth. In their growth model, where data appears as one of the factors of production, nonrivalry of data gives rise to increasing returns to scale when data is combined with other inputs. The intuition is that each unit of data can be used by all units of other inputs simultaneously. A larger stock of complementary labor or capital allows each unit of data to be better exploited, raising the average product of data. An implication is that access to the same nonrival data results in larger firms with more complementary inputs being more productive than those with fewer inputs. This will tend to increase average firm size in the economy and can potentially stifle competition by representing a barrier to entry for smaller, data-poor firms.

Crucially, the increasing returns to scale extend as far as access to nonrival data is granted. The more data is created and made available to all firms in the economy, the more goods or services can be produced with the same amount of other inputs. Accumulating more data thus resembles a process of technological change. Given the inherently global reach of information networks, wide access to nonrival data could enable increasing returns to scale at the national or even the international level. Like ideas in Romer's (1990) endogenous growth model, the nonrivalry of data can thus give rise to sustained growth through increasing returns to scale.

This conclusion is not shared by Farboodi and Veldkamp (2019), who emphasize data's role in reducing uncertainty. Because uncertainty is bounded by zero, there are limits to the efficiency gains that data can deliver. In their setup, data accumulation resembles a process of capital accumulation, rather

than a process of technological change. The result is that data accumulation is subject to decreasing returns and thus cannot deliver a boost to long-term economic growth.

Equity

We have described how data can be a source of economic efficiency and growth. Three types of agents play a role in generating value with data, and thus may have a legitimate claim on the returns it generates: (1) the data subject—who the data is about, and has some preference over keeping it private[3]; (2) the data collector—who pays the cost of recording and protecting the data; and (3) the data processor—who performs costly analysis on a dataset to extract insights and knowledge. Each of these agents has their own interests and preferences, and their decisions over data collection, processing, transferring, and sharing will have an incidence on the others.

How will the returns to data accrue across these agents? A key question is to what extent the value of data comes from each individual data point, or from its agglomeration and subsequent analysis. The answer is likely to differ according to the type of data and the context, but an important factor is the degree of substitutability between data and other factors of production. For instance, some types of data may require very advanced and proprietary analytical tools to convert into useful information. In the case of big data sets used to train machine learning algorithms, most of the value likely comes from the analysis provided by the high-skilled labor required to process the data. In other cases, information may be extracted with less analysis, as when individual data is used to provide a targeted product. The value of such data may differ greatly depending on characteristics of the data subject: being able to target an ad to a high-net-worth individual may offer a much higher return to an advertiser than targeting to an unemployed worker. Here the data collector may be able to extract most of the rents from the data.

Implications of data proliferation for inequality will depend on the market power enjoyed by data subjects, collectors, and processors. If a data collector enjoys market power, then obtaining granular information on their clients may enable them to extract considerable rents through the implementation of price discrimination strategies. Data may also be a source of market power if a stockpile of data acts as a barrier to entry that deters competition. We discussed some of the arguments that might lead to market concentration

[3]Some data involves multiple data subjects. For instance, one person may divulge that they are friends with another or that they both attended a political gathering. If the first decides to share this data, it may have implications for the second.

in data in the previous section, which include network externalities and learning-by-doing feedback loops.

In the context of financial services, the proliferation of data can alleviate adverse selection problems that exclude vulnerable populations from credit markets—a problem that is particularly acute in emerging and developing economies. One means of doing so is to facilitate the sharing of individual financial data across financial service providers, including through the adoption of open banking (Box 1). Indeed, there is empirical evidence that banning information sharing through regulation aggravates adverse selection problems and leads to more credit rationing. Liberman and others (2018) exploit a one-off legal reform in Chile that called for the deletion of default information from the national credit bureau, estimating that the measure had the effect of excluding vulnerable borrowers from obtaining future loans.

Whereas data on borrowers' financial history has been collected, processed, and shared for decades, a potential source of further reductions in information asymmetries is the use of nontraditional data in finance. For instance, information collected in the context of online services, including social habits, payment of utility bills, and other traces of economic and social activity, may form the basis for evaluating the creditworthiness of a borrower who has not had previous interactions with a financial service provider.[4] This may alleviate adverse selection effects, address collateral constraints, and broaden the number of clients able to obtain a loan, but it may also lead to the exclusion of those that exhibit traits associated with risky financial behavior.

More broadly, data-driven assessments raise concerns about introducing bias into credit decisions that may reflect average outcomes, but that are inconsistent with social norms and values. Consider a pattern of lending that effectively reduces access to individuals based on location or race. The use of artificial intelligence algorithms—which produce accurate predictions but often lack a structural interpretation—may be perceived as a discriminatory black box that loan officers will not be able to explain to their customers or to regulators.

Importantly, the efficiency gains from improved data availability in finance may be unevenly shared across borrowing firms within an industry. Begenau, Farboodi, and Veldkamp (2018) argue that the emergence of big data in the financial sector benefits larger incumbent firms more than small firms, who produce less data. This differential impact on the cost of finance contributes to increasing returns to scale as larger firms generate more data, lowering their cost to finance expansion, which in turns generates more data. The

[4]Bank for International Settlements (2019) discusses the use of data held by nonbank Big Tech firms to provide financial services such as credit scoring and loan monitoring.

implication for market structure is that data-based lending will tend to favor more concentration in production.

Finally, the availability of granular data can undermine the risk-sharing function of insurance. For instance, knowledge of a preexisting medical condition may lead an insurer to charge higher premiums or to deny coverage (Arrow 1963). The proliferation of data potentially enables the insurance company to discriminate based on a wider range of criteria, such as the past demonstration of risky driving habits or social connections to people that lead unhealthy lifestyles. Such discriminatory practices will tend to shift welfare toward those who hold more information and can lead to the exclusion of vulnerable individuals from insurance markets.

In an extreme, perfect information derived from data acquisition and processing may undermine the basic risk-pooling function of insurance markets. As in the case of bans on the use of data about preexisting conditions in many jurisdictions, policies must be designed to limit the factors on which the conditions of coverage can be determined.

Stability

The proliferation of data has implications for stability in sectors where cyber-attacks can undermine public trust in service provision (such as financial services) or can directly compromise the sector's operations (such as power grids). Implications for financial stability are thought to work through three main channels.[5]

First, data is likely to have implications for market structure among banks and among borrowing firms that affect risk-taking behavior and the well-studied concentration-stability trade off (Dell'Ariccia and Marquez 2004; Vives 2016). Particularly where an incumbent lender can hoard their customer's financial data, adverse selection problems may impede competition from other lenders to whom the customer is unknown. This can make the market less contestable and increase concentration in lending, and allow the lender to charge high interest rates.

Second, as financial firms rely increasingly on data processing in their business, customer data becomes vulnerable to cybersecurity risks, which threaten public trust in financial institutions. As the provision of financial services becomes increasingly reliant on data, the system's stability depends on financial firms and their service providers adequately managing cybersecurity risks

[5]Financial Stability Board (2017) flags several channels for data-driven financial technology to impact stability, including some of those emphasized here.

to avoid wide-scale disruption. Indeed, the protection of user data represents an important cost for financial institutions who are interested in building and preserving their reputation as reliable custodians of customer data. As discussed in the section titled "The Building Blocks of a Market for Data," the private incentives for investment in cybersecurity may be smaller than the social optimal.

Third, data proliferation in finance influences operational risks facing financial institutions. These risks may become systemic from the use of common interconnected systems, including third-party cloud computing services, which create nodes of risk. However, the use of more sophisticated cloud-based systems may represent a substantial improvement in operating standards for institutions with lower capacity. In addition, nonrivalry of data and the use of interconnected systems allow for the geographic diversification of operational risks.

Box 1. Data Sharing in Banking: From Credit Bureaus to Open Banking

As discussed in the section titled "The Building Blocks of a Market for Data," credit bureaus have long served to reduce information asymmetries in finance through reciprocal data sharing. As the cost of collecting data on customer interactions falls, might there be a case for broadening the types of data that is shared among banks? Many jurisdictions are looking to spur innovation and competition in the financial sector by adopting open banking principles. This approach typically involves regulation that requires banks to share all their information about a customer with a designated third party—including competing banks or other licensed financial service providers—at the customer's request. Although the details of what data is to be shared vary across jurisdictions, these usually include granular data on transactions and account balances, and in some cases extends to mortgage or consumer loan balances and payments.[1]

It is argued that open banking reduces barriers to entry by breaking incumbent banks' monopoly over their customers' financial data. At the same time, new fintech entrants are absorbing new sources of information on consumers to provide additional financial services, often in competition with incumbent intermediaries. Questions arise on the perimeter of entities that should be subject to data sharing requirements to engender a level playing field. For example, should there be reciprocity in data sharing between technology providers and banks, to the extent that they compete in the provision of similar financial services?

This again raises the question of whether the value of data stems from the individual data points, or from the agglomeration and analysis of big data. In the case of financial services, both are likely to be important. Consider a new bank that is considering whether to extend a loan to a new client and on what terms to do so. In the absence of historical data on the client, the information asymmetries involved will be very large. If the customer were able to provide their comprehensive financial records to the new bank—say, by invoking their right to data portability under open banking regulations—the new bank would have some information on which to base their assessment, lowering the information asymmetry they face. However, they would still be at a disadvantage with respect to a large incumbent bank, who would also have access to a large database against which to compare the individual's data. As access to big data allows lenders to improve the precision of their creditworthiness assessment through analysis, it will continue to act as an advantage to incumbents. Should incumbent banks then be compelled to share their entire customer databases on an anonymized basis, on procompetition grounds?

[1]Beyond data sharing, open banking has typically involved allowing access to direct payment initiation, but the implications of this dimension are outside the scope of this paper.

This page intentionally left blank

CHAPTER

5 Data Policy Frameworks

We have argued that the use of data in the economy holds the promise of substantial efficiency gains but can also potentially lead to market concentration and a stifling of competition. Likewise, externalities arise with regards to individual privacy and the need to ensure that private decisions to protect data underpin trust and stability. Data policy frameworks must balance trade-offs across many competing objectives, and doing so requires an integrated approach.

Interventions that address a single facet of data are likely to generate suboptimal outcomes. For instance, a push to tighten privacy regulations may be effective at protecting consumer rights but may generate unforeseen harm to efficiency and competition. Likewise, granting data collectors extensive rights to collect and process data as they see fit is likely to create strong incentives for data generation—and commercial benefits for the data collector—but may lead to data hoarding that precludes the broader efficiency gains that could be achieved through wider access. It may also lead to disregard for user privacy and for stability.

In practice, current data policy frameworks tend to involve individual laws and regulations that are targeted at specific aspects. In this section, we discuss how data policy frameworks should be reformed to address four broad concerns that are arising with the status quo: market opacity, concentration and market power, financial instability, and international fragmentation.

Market Opacity

Whose data is it, when can it be collected, what can it be used for, and who will have access to it? Participants in the data economy may not have a clear sense of the answers to these questions, which are key aspects of the transac-

tions they are continually engaged in. How rights and obligations over data are defined and allocated will determine many of the implications of data in the economy. Because control and access rights over data are often unclear, data is likely being overused, with privacy insufficiently respected in the absence of meaningful consent. Zuboff (2019) describes contemporary online data markets as "one-way mirrors," with data collectors designing platforms such that users have only a vague sense of the personal data that is being collected and are entirely unaware of what is done with the data after it has been collected, including being sold to third parties.

For data markets to operate efficiently—balancing the efficiency gains from data processing against the data subjects' preference for privacy—all parties must understand the terms of the economic decisions that are taking place. As a first step, efficient data markets require clarity as to rights and obligations. The Coase (1960) theorem states that, as long as the property rights over the goods being allocated are well defined and respected, then a competitive market will be able to achieve a social optimum. An additional aspect of Coase is that the initial allocation of property rights across agents will not determine the final market allocation. However, the literature has tended to argue that the theorem is unlikely to apply to data (Acquisti, Taylor, and Wagman 2016). In the presence of market power and asymmetric information—most users are not aware of how their data are being used—the allocation of data is likely to depend on the ownership and access rights assigned to data subjects, collectors, and processors.

But even with clear rights, the mechanisms by which users engage in data transactions should clarify what is happening. For instance, will the data collector use the data for a specific purpose, or will they reserve the right to use it again later, including by selling it to a third party? As we discussed, nonrivalry makes data a potential source of economies of scope for the data collector. There will thus be a tendency for data collectors that have acquired data for one purpose to use it for many other purposes as well. This usage of data can generate efficiency gains but may impose new costs on the data subject in the case of personal data. For the decision to be socially efficient, the data subject should have control over each usage of their data. The European Union's General Data Protection Regulation recognizes and seeks to address many of these concerns (Box 2) by including a special policy framework for personal data that grants control rights to the data subject.

Consent is the mechanism by which control is exercised, but can meaningful consent be operationalized? Users are often made to agree to cumbersome legal contracts in exchange for the use of relatively minor services, and most accede to the terms and conditions without being aware of what they are agreeing to. In many cases, the data exchange is designed to appear

like a minor side-product of the main transaction—not worthy of much engagement—whereas it is often the most economically significant. The more complex are the terms of the transaction, the less likely that meaningful consent will be achieved, with optimal allocations remaining elusive.

Trust and reputation play an important role in data transactions, because even the perception that a data collector misused data about their user can lead to a withdrawal of their participation in the future. Terms must be presented in simple language that users will understand and should involve simple transactions for a limited use and period. Irreversible decisions about data may also be problematic, because they require users to evaluate costs that may materialize in the distant future. The right to be forgotten requires that a data collector or processor delete old data about a data subject at their request. This makes data transactions more salient and transparent to evaluate, because they involve an exchange for a well-defined duration.

To whom data rights and obligations are assigned will also have important implications for the macroeconomy. Acquisti, Taylor, and Wagman (2016) emphasize that the allocation of data access rights will have a bearing on how rents are distributed across data subjects, collectors and processors. Arrieta-Ibarra and others (2018) argue that the current approach of assigning broad rights to data collectors and processors essentially treats data as capital. A mix of low market power, ignorance, nontransparent data practices, and historical path dependence have allowed data processors to extract the rents of data subjects. They call for data governance frameworks to treat data as labor, whereby data processors would pay data subjects for the right to process their data.

Jones and Tonetti (2018) provide a strong argument that the implications of data for growth and equity will depend greatly on how policy frameworks address the question of control over data. In their model, they study the welfare implications of alternative data policy regimes, including those where users own data, where firms own data, and where data is made publicly available to all. When firms are granted full control of the data they collect about their customers, they tend to hoard it rather than share it with current or prospective competitors. This is bad for efficiency for two reasons. First, the firms limit the economies of scale generated by nonrivalry to their inputs, precluding its use by factors in other firms and erecting a barrier to entry that stifles competition. Second, firms do not fully consider their users' preferences for privacy, and thus tend to collect too much data from data subjects. When the policy framework grants firms ownership and control over data, this leads to too much data being collected from users, but not enough sharing of data across firms.

In contrast, when data subjects are granted ownership of the data that are collected about them, they are able to trade-off their preference for privacy with their desire to receive better data-fueled consumption, which tends to limit the amount of data the firm collects from them. But on the other hand, they ignore the firm's desire for market power and profits, and thus tend to allow wider access to their data once it has been collected. When the policy framework grants users ownership and control over data, the amount of data collected and the extent to which it is shared are much closer to the social optimum.

Concentration and Market Power

Crémer, de Montjoye, and Schweitzer (2019) argue that access to data is an important competitive parameter in the modern economy. Do current market dynamics suggest that data access is being granted in a way that ensures adequate competition? Incumbents in the data economy appear to be earning large rents, as reflected by high reported profits and equity market valuations, and many digital markets currently feature high degrees of concentration (Furman 2019). This may reflect a practice of hoarding data on their customers, creating a barrier to entry that is stifling competition from smaller firms in some cases.

Policy tools to increase data access include requiring *portability* of individual user data. This corresponds to granting users the right to access and transfer the personal data that collectors and processors are holding about them. One aim of portability requirements is to promote competition by lowering the data subject's cost of switching to a competing service or of multihoming across multiple services.[1] Implementing portability imposes costs on data processors, who must build an interface from which users can access their data. In a highly concentrated market, this may have a meaningful impact because competitors will obtain user data in a single format: the format used by the incumbent. However, if the market features several incumbents with different data formats, the lack of a common standard may make portability too costly for competitors to implement, thus reducing its impact on competition. In studying the European context, Crémer, de Montjoye, and Schweitzer (2019) propose that data portability could be imposed on specific dominant firms, where lock-in effects may be particularly pronounced.

A complementary policy measure to portability requirements is to require *interoperability* of data across platforms through common standards. This has tended to be implemented through sector-specific regulations, where

[1]Another important objective of portability is to reduce opacity by revealing to users the data that is held about them.

the type of data involved can be well defined. An example is open banking frameworks—deployed by Australia, the United Kingdom, and European Union, among other jurisdictions—where portability and interoperability are mandated together to encourage competition among banks and nonbank providers of payment, saving, and lending services (see Box 1).

An important feature of many current data markets is that they occur on multisided platforms that can solve coordination problems between agents that do not interact directly (Evans and Schmalensee 2014; Rochet and Tirole 2006). Consider a large social media company that offers its platform at no direct charge to users but collects data on user activity that it sells to advertisers at a price well above cost. Although the direct price paid for the user's data is zero, the data generates sizable rents to the processor through their other commercial activities. The normative question arises as to what the distribution of these rents ought to be, and whether the provision of the social network's services to the user represents appropriate compensation (Arrieta-Ibarra and others 2018).

A recent report prepared for the UK government by Furman (2019) considers policy responses to concentration in two-sided data markets. They argue that policy should facilitate sharing of personal data across platforms by mandating portability and interoperability to operationalize data access. Such an approach is seen as a means of bolstering competition and the contestability of markets. The notion would be that the threat that a user could switch to another service would discipline the incumbent platform and lead them to share a larger portion of the rent.

However, a limitation to the idea of data portability as a solution to these multisided market problems is that the underlying cost structure may involve strong economies of scale, such that the market tends toward concentration, leaving the user with limited options. And network externalities may make switching to another platform unattractive unless they have enough scale, such that the option of portability may never be valuable enough to exercise. The extent to which interoperability standards would mitigate these forces is open to discussion. More broadly, Evans and Schmalansee (2013) point to the limitations of applying antitrust theory developed for single-sided firms to multisided platforms, suggesting more work is needed including considering how to regulate market structure in such cases, including arising from the access to and control of data.

Beyond clarifying the legal framework for data ownership and access, is there a case for alternative governance frameworks for data, such as public data trusts or data utilities? Furman (2019) advocates for the creation of a pro-competition data markets unit that could compel firms to share their data with other data processors in particular cases. Another model is to provide

some types of data that are particularly valuable and sensitive through a *public utility*, designed to provide socially optimal levels of privacy and cybersecurity, while providing access to exploit strong returns to scale and promote competition. This reflects that certain economic characteristics of data make it resemble a *public good*: it is nonrival and only partially excludable. A further dimension is that it may display strong enough economies of scale and scope to generate a *natural monopoly*, which may lead to data being insufficiently shared by a profit-maximizing data collector.

To manage the trade-offs involved across access, privacy, and cybersecurity decisions according to social preferences, access to these data could be handled by the state or by a private company subject to strict oversight. Indeed, official economic statistics have been treated as public goods and managed by government agencies for many years. More recently, the proliferation of biometric identification has allowed for the proliferation of analysis using large administrative data sets handled by public institutions. However, recent cyber-attacks on large public data sets have exposed the personal information of millions of citizens and highlighted the limits of this approach. Indeed, no one model is likely to offer a panacea, and approaches will result in different policy preferences across objectives.

Financial Instability

As discussed previously, the proliferation of data in the financial sector has implications for systemic stability through multiple channels. It is important to note that not all these channels necessarily imply an increase in the level of systemic risk, particularly if the risks are managed adequately by sound policies. However, current approaches to data in finance may be allowing a buildup of certain types of stability risks, whose mitigation warrants prudential policy responses.

First, large-scale data breaches at important banks and data intermediaries underscore the need to ensure that financial institutions spend enough on securing client data. Indeed, Kashyap and Wetherilt (2019) argue that the private incentives for investment in prevention and recovery capabilities may be insufficient to adequately mitigate these risks. Private investment in cybersecurity is subject to an externality, because a security breach at one institution can disrupt other firms directly or indirectly, by undermining public trust in the data security of the broader financial system. How to implement mitigating policy measures in practice is subject to severe challenges, including the development of monitoring and surveillance systems, and redesign of supervisory and enforcement tools.

Second, the potential systemic financial stability risk implications of the rapid increase in the use of third-party cloud service providers—that are growing as critical repositories for personal and commercial financial data collected by financial intermediaries—remains an important issue for consideration by financial supervisors, as flagged by the Financial Stability Board (2017, 2019). The Financial Stability Board notes that usage of such services, while still at an early stage, may reduce operational risk at the individual firm level—for example, by increasing cyber resilience and supporting business continuity—but could "also pose new risks and challenges for the financial system as a whole, particularly if risks are not appropriately managed at the firm level, and if the complexities and interconnectedness of third parties and their usage continue to grow" (Financial Stability Board 2019). In particular, given that the cloud services market is highly concentrated, as data and core functionality migrate to the cloud, single point of failure risks could emerge. Nonetheless, it is an open question whether such risks are materially different than those posed by existing data centers and services. Developing policy approaches to managing potential emergent risks faces several challenges, especially with regards to domestic and international coordination. Indeed, there may be a need for both greater international coordination across financial authorities, and also between financial authorities and their counterparts responsible for information technology safety and security. A further challenge for financial authorities is to deepen the scarce skills needed for effective supervision of technology providers, which lie at the intersection of economics, finance, and computer science.

Lastly, data-driven credit analytics are powering the growth of financial services provision and increased inclusion in many regions (Sahay, von Allmen, and Lahreche, forthcoming). As noted previously, this has many benefits including by alleviating credit constraints for high return projects, including in emerging market and developing economies. However, given the very recent development of such models and the fact that the data used do not in many cases span a full financial cycle, questions arise as to the resilience of this new lending were economic and financial conditions to deteriorate (Claessens and others 2018). Moreover, data-driven lending provided by online platforms has also raised important consumer protection risks, including those that arise from person-to-person lending platforms.

International Fragmentation

As countries adopt very different approaches to data policy frameworks, there is a risk of international fragmentation in data and goods trade. This would preclude important potential gains from cross-border activities. Because data is nonrival and can be transferred anywhere at virtually zero cost, it is the

ultimate mobile factor of production. Just as nonrival data can give rise to sustained growth through increasing returns to scale within the firm or across firms, scale can also be achieved across countries. But will it be allowed?

Global real and financial integration have been matched by global integration of data and information flows via the worldwide web and other networked information-sharing technologies within and across individuals, firms, and governments. Reflecting in part the decentralized origins of the internet, its global governance has largely been reflected in a distributed approach to agreeing on protocols and standards (for example, FTP for file sharing, SMTP for mail services, IP for internet addresses, etc.) with some degree of coordination offered by nongovernmental bodies (for example, ICAN).

Although there is no global governance framework for data per se—though agreements on trade in services tend to have elements focused on allowing free flow of data—approaches to data rights vary greatly across countries, across sectors, and even across subnational jurisdictions. Indeed, many national approaches have developed to address consumer protection and financial stability concerns, namely to maintain privacy of data subjects and offer some protections in case of theft of data or identity. In some cases, countries have opted to impose some degree of data autarky by passing laws that require data on their citizens to be stored within national borders (Box 3).

National approaches to data governance have important international ramifications, because cross-border provision of data services must comply with local frameworks. No sooner had the General Data Protection Regulation come into force did companies around the world come to realize that its rules would apply to them to the extent that some of their individual customers reside in the European Union. In other words, the General Data Protection Regulation generated sizable policy spillovers. Other jurisdictions are also taking note and considering whether the European approach will be right for them. Moreover, instability and loss of trust in one jurisdiction could impose costs on financial systems in other (interconnected) jurisdictions. As such, individual jurisdictions may not fully internalize the externalities they impose on the global system if their own regulation of data and privacy is focused exclusively on domestic considerations.

Although we shouldn't expect all countries to handle issues of innovation, privacy, and security the same way, there is a strong case for international dialogue and cooperation to ensure that the digital economy does not become subject to undue fragmentation. Aspiring to the best principles of privacy and individual rights should ideally not set off a global scramble to fragmented policy approaches leading to localized data markets that could undermine the many potential benefits of cross-border data sharing.

There is a need for coordination to develop international frameworks that set minimum standards balancing interests of growth and competition with national and individual privacy concerns.[2] Of course, one must recognize the challenges involved in developing such a framework, given the diversity of current national approaches and legitimate concerns over sovereignty, privacy, and national security.

[2]Although very early to judge, elements to explore in such frameworks could include modalities for allowing efficient management of data across national boundaries while addressing concerns, potentially utilizing the services of trust-generating entities—which may include new private-public partnerships or even certain financial centers—that store and validate data subject to standards that provide assurance to national authorities.

Box 2. The European Union's General Data Protection Regulation

The General Data Protection Regulation (GDPR) implemented by the European Union in May 2018 has set the initial global standard for modernizing data policy frameworks by defining, clarifying, and protecting the rights of European Union residents over their personal data. Noncompliance of these data rights and obligations exposes data processing firms to large fines, regardless of their country of origin. Given the European Union's size and interconnectedness in the global economy, the implications of GDPR extend across international borders.

A key feature of the GDPR is its establishment of a framework that specifies the rights of individuals who are the subject of data, and the obligations of the companies that collect, store, and analyze it. Under the regulation's principle of data minimization, data handling should involve only as much personal data as required to accomplish a lawful purpose. Data collected for such purposes in principle is not to be repurposed without further user consent. The GDPR also gives individuals new, expanded rights over their data. These include a "right to be forgotten," a "right to erasure" to have personal data deleted, a "right to rectification" for information to be corrected, and a "right to portability" to retrieve or transfer one's personal data in an electronic format at no charge. The regulation also includes provisions aimed at ensuring that adequate resources are destined to securing personal data. Specifically, data collectors are required to anonymize user data they store—for instance, through encryption or tokenization—so that any given piece of information cannot be readily matched to an identity.

The development of GDPR as a new policy framework has important economic consequences and could help engender a better-ordered market for data that balances opportunities and risks by clarifying how externalities from privacy and excludability are to be handled. Some observers have noted that an emphasis on privacy protection may impact innovation by acting as a tax on digital technologies or that the compliance costs for start-ups may be very high, potentially reducing competition, which could have adverse impact on consumers. There is some early evidence that GDPR has had an effect on the ability of e-commerce firms to attract users and generate revenue (Goldberg, Johnson, and Shriver 2019) and to raise funding (Jia, Jin, and Wagman 2018). Further study is needed on the broader economic implications and trade-offs associated with policies that seek to strengthen consumer privacy.

Box 3. Data Localization

Many multinational companies store data on their customers on global networks that span national borders. Some governments have introduced data localization laws that either require or encourage companies to store individual data on their citizens within national borders, or even restrict the transfer of individual data across borders. In some cases, restrictions are limited to data from specific sectors that are deemed particularly sensitive to personal or national security (such as health or finance), but in others the requirements are broader. Data localization laws raise several conceptual issues.

Protectionism: Data autarky may stifle competition and growth in the digital economy. By raising the cost of collecting, transferring, and storing data across borders, data localization (including local storage requirements and restrictions on cross-border data flows) may act as a trade barrier to protect local firms from competition with foreign incumbents, with the unintended consequence of reducing innovation and integration. Localization is expensive for data-intensive companies to implement, as it requires the installation of data infrastructure in each country whose users it serves. Although initial costs are likely to be borne by large incumbent firms, implications could be more widespread. For instance, if imposed by smaller jurisdictions, the cost of compliance may prove high and lead providers to curtail services.

Sovereignty: Localization requirements give national governments the ability to exert sovereignty over their citizens' personal data. In the face of concerns over covert surveillance by foreign intelligence agencies, some jurisdictions have invoked an interest in protecting their citizens' right to privacy through more stringent data law and regulation. Data localization has been a part of these discussions, as it, in principle, allows domestic governments more control over the personal data stored on corporate servers. Whether localization laws end up leading to greater privacy protection ultimately depends on the relative protections offered by the domestic and foreign jurisdictions.

Cyber risk: Data localization laws may mitigate or amplify cyber and national security risks by decentralizing the storage of data across countries. Although this expands the number of potential targets for cyber-attacks—increasing the cost of protecting the whole network—it also implies that any single data breach may have a smaller chance of causing globally systemic damage. When the information relates to core infrastructures such as financial services, health provision, or energy distribution, that are all relevant for national security, local data storage may also make it easier for national governments to physically protect these infrastructures. Whether risks can be mitigated by local oversight will depend crucially on scale and capacity.

This page intentionally left blank

CHAPTER

6 The Case for an Integrated Approach

This paper has sought to integrate different strands of literature in economics and finance that provide building blocks for understanding the economic implications of data. These include how data is produced, why it is demanded, and how the market for data can develop with different organizational structures. We have emphasized three economic characteristics of data that can have implications for markets: data is nonrival, involves privacy externalities, and is only partially excludable. Whereas each of these characteristics is not unique to data—allowing us to draw on extensive literatures that consider each in other contexts—they certainly set the economics of data apart from the economics of oil.

We view that data policy frameworks have first-order implications for macroeconomic growth, equity, and stability. An approach in which a sectoral regulator considers a privacy, innovation, competition, or financial stability objective, in isolation, may well have implications for other objectives, with a risk of individual policy actions being at cross-purposes when considered as part of a whole package. Balancing these competing objectives in our view calls for an integrated approach to data policy frameworks that involves coordination among many national organizations, including central banks, ministries of finance and economics, financial regulators, consumer protection agencies, privacy regulators, and competition agencies.

Given the heterogeneity of data varieties—including whether the data describe sensitive personal or commercial traits—there is a case for sector-specific treatments and exceptions. However, an integrated approach should be mindful to ensure that these do not generate unintended trade-offs across the economy. For example, efforts to boost competition by mandating data interoperability in finance through open banking regulations can create an unlevel competitive playing field between regulated incumbents and Big Tech challengers.

How should data policy be designed to protect privacy while also facilitating wide dissemination of data—a social good—and preserving as much competition as feasible? The paper poses many questions to which there do not appear to be easy answers. We argue that four concerns about the status quo merit changes to data policy.

First, rights and obligations over data must be clarified for the market to function efficiently. An early insight from the literature on the economics of data is that the implications of data for growth and equity will be determined by who controls access to data in the economy and in providing clarity on the distribution of economic returns from use of this data.

Second, large incumbents in the data economy appear to have gained substantial market power based on a strategy of hoarding customer data. This calls for policies that can encourage user control over data and complete more competitive markets, and some have called for mandates on data sharing across firms to boost competition.

Third, the proliferation of data may not be being met with sufficient investment in cybersecurity, and this may be reducing the stability of the financial system. Other stability concerns arise from the over-the-cycle resilience of data-driven credit provision and use of third-party service providers to house critical data infrastructures.

Fourth, there is a risk of international fragmentation in data markets, which could reduce the potentially sizable efficiency gains from the economies of scale and scope inherent in free movement of data across borders while also reducing financial resilience arising from the distribution of information on large, secure global networks. Many countries are moving fast to modernize national data policy frameworks with varying degrees of concern placed on data privacy, consumer protection, national security, and competition. So far, there are limited international discussions to foster international cooperation on data that span all major data economies. This has prompted calls for an international framework or standards for the management of data or even for a centralized global body to set standards (Tett 2019).

These and other questions deserve further thought, and the analytical ingredients synthesized in this paper could be used as a basis to develop an assessment of policy relevant trade-offs.