

The Social Epistemology of Bureaucracies

PHILIPP KREMERS

Department of Philosophy, Georgetown University, Washington, D.C., 20057, USA

Abstract: We admire some bureaucracies for their ability to efficiently organize the production of knowledge on a large scale. At the same time, bureaucracies are commonly associated with serious epistemic shortcomings. As Max Weber noted, they have a reputation for turning their members into ‘specialists without mind.’ This article resolves this apparent tension by advancing a systems-oriented social epistemology of bureaucracies. I show that bureaucracies indeed have the potential to produce and distribute knowledge in a fast and reliable way. However, this epistemic advantage only manifests under certain ideal conditions (e.g., when general conflicts between assigned social roles and personal interests are absent, when cost-efficient communication technology is available, and when the bureaucratic social system is located in a relatively stable environment). Furthermore, I show that the fragmentation of knowledge within bureaucracies facilitates the disintegration of understanding, which is a serious epistemic downside that emerges even under favorable conditions.

“Comprehension is not a requisite of cooperation.”

—Councillor West, *The Matrix Reloaded*

§ 1. Introduction

The term ‘bureaucracy’ has carried pejorative connotations ever since its invention in the mid-18th century. Nevertheless, Max Weber, arguably the most influential theorist of bureaucracy, had high hopes for these social systems.¹ Weber argued that the distinctive role of trained specialists within bureaucracies makes these social systems superior on epistemic grounds: “bureaucracy is superior in knowledge” (Weber 1964 [1921], 339). There are, indeed, many bureaucracies that perform remarkably well epistemically. Consider the following example:

SUCCESS STORY: Gaspard de Prony, a French mathematician, was tasked shortly after the French Revolution with producing elaborate tables of decimal approximations of logarithmic and trigonometric functions for the newly established Bureau de Cadastre. These tables were essential both for naval navigation as well as the mapping of territory. However, de Prony was aware that even skilled mathematicians struggle to compute decimal approximations of logarithmic and trigonometric functions without error. Fortunately, he stumbled upon a copy of Adam Smith’s *The Wealth of Nations* in a bookshop in London. In this book, Smith famously points out that the division of labor (e.g., in a pin workshop) can dramatically enhance the quantity and quality of production. Inspired by Smith’s point, de Prony realized that one “could manufacture logarithms as easily as one manufactures pins” (Grier 2005, 36). He broke

¹ On the history of the concept of bureaucracy, see Joukovskaia (2023). As Höpfl (2006, 13) observes, it is peculiar that Weber chose the pejorative term ‘bureaucracy’ despite his generally positive assessment of these social systems.

down the task of computing logarithmic and trigonometric tables into numerous simple tasks and distributed them among a large group of workers—many of them without any noteworthy mathematical training. His approach was crowned by success: “the trigonometric tables surpassed all previous ones in accuracy, and calculations of logarithms wholly from scratch had not been undertaken since the seventeenth century” (Daston 1994: 189).²

This example illustrates the enormous epistemic potential of bureaucracies. The bureaucratic structure of de Prony’s workshop allowed him to solve a complex epistemic problem in a remarkably reliable and fast way. Later, de Prony’s method even inspired Charles Babbage to invent the Difference Engine, the world’s very first mechanical computer (Daston 1994, 195 ff.). Thus, even though the connection is admittedly quite indirect, the epistemic benefits of modern computing technology can be traced back, at least in part, to the bureaucratic method.

At the same time, it is well known that bureaucracies do not always lead to favorable epistemic outcomes. Consider the following example:

FIASCO: On the morning of December 7, 1941, the Japanese Empire launched a surprise attack on a US naval base at Pearl Harbor in Honolulu, HI. That morning, radar operators of the US Army detected a large group of approaching aircraft. When they reported this to their commanding officer Lieutenant Kermit Tyler, he assumed that the radar blip was just a group of expected B-17 bombers and dismissed the report. Tyler and his subordinate radar operators lacked access to contextual knowledge that might have allowed them to correctly interpret the signal. Meanwhile, higher-ranking officers who did have this contextual knowledge were unaware of the radar signal itself (Downs 1966, 91–93).

This example vividly demonstrates the epistemic downsides of bureaucracies. Because knowledge was highly fragmented within the military bureaucracy, no individual soldier ended up having a proper understanding of the situation. As a result, the attack caught the soldiers off guard and could inflict devastating damage on the naval base.

How is it that some bureaucracies give rise to strikingly good epistemic outcomes in cases like *SUCCESS STORY*, while they lead to severe epistemic shortcomings in cases like *FIASCO*? This is the question that I want to answer in this article. I will first outline an ideal social epistemology of bureaucracies that is inspired by Weber. In short, this ideal theory holds that bureaucratic social systems are overall epistemically good because they allow their members to produce and distribute knowledge in a fast and reliable way. We will see that this ideal theory can explain what makes de Prony’s bureaucracy in *SUCCESS STORY* epistemically good, but that it fails to make sense of the epistemic shortcomings in *FIASCO*. I argue that this ideal social epistemology of bureaucracies relies on a superficial characterization of the human and non-human components of these social systems, their environment, as well as the relevant epistemic goods. Thus, we will see that bureaucracies only promise to give rise to good epistemic outcomes if general conflicts between assigned social roles and personal interests are absent, if cost-efficient communication technology is available, and if the bureaucracy is located in a relatively stable environment. Furthermore, it will turn out that the fragmentation of knowledge within

² For a more detailed account of this case, see Simon and Newell (1958), Grattan-Guinness (1990), Daston (1994), and Grier (2005, ch. 2). The significance of this case for the epistemology of testimony has been discussed by Shieber (2013).

bureaucracies facilitates the disintegration of understanding. This is a general epistemic downside of bureaucracies that emerges even under favorable conditions. We will see that this non-ideal social epistemology of bureaucracies can explain both the epistemic shortcomings in *FIASCO* and the good epistemic outcomes in *SUCCESS STORY*.

This non-ideal social epistemology of bureaucracies is not only relevant because it helps us to better understand these social systems in themselves, but it also allows us to identify countermeasures that can help us to mitigate their epistemic downsides. Furthermore, if it turns out to be true that there is a *general* tendency towards bureaucratization in modern society, as some have argued (Parkinson 1958; DiMaggio & Powell 1983; Weber 1993 [1930]; Graeber 2005, 149–153), then this outcome will turn out to be significant for the epistemic assessment of virtually any domain of society.

I will proceed as follows. First, I will clarify what a bureaucracy is (§ 2). Second, I will elaborate on some methodological complications for any attempt to evaluate bureaucracies from an epistemic point of view (§ 3). Third, I will put forward a Weber-inspired ideal systems-oriented social epistemology of bureaucracies (§ 4). Fourth, I will expose the blind spots of this ideal theory and propose a non-ideal theory of bureaucracies that is sensitive to their human and non-human components, their environment, and all epistemic goods that are relevant for their epistemic evaluation (§ 5).

§ 2. Understanding Bureaucratic Social Systems

In this section, I outline the epistemically relevant features that bureaucracies paradigmatically exhibit.

Weber's model of bureaucracy has attained canonical status in the social sciences. He famously refrains from defining bureaucracy in terms of necessary and sufficient conditions but, rather, sketches an 'ideal type.' Thus, even if a social system does not perfectly embody this ideal type of a bureaucracy, it may still count as one. The most prominent articulation of Weber's ideal type of bureaucracy can be found in his unfinished monograph *Economy and Society*, which was published posthumously in 1921/22 by his wife Marianne Weber.³ His most detailed discussion appears in an unfinished chapter titled 'Nature, Conditions, and Development of Bureaucratic Domination [*Herrschaft*].'⁴ There, Weber argues that bureaucracies paradigmatically exhibit the following six properties:

- 1) Weber claims that a bureaucracy "is based on the general principle of precisely defined and organized across-the-board *competencies*" (Weber 2015 [1921], 76). Thus, he takes for granted that a bureaucracy engages in a shared task and tries to accomplish this shared task by means of a "rigid division of labor" (Weber 2015 [1921], 76). More specifically, according to Weber, a bureaucracy consists of a number of offices [*Ämter*], which are occupied by officers [*Beamte*]. These officers are full-time employees that exert authority in virtue of their position within the

³ I mostly rely on Tony and Dagmar Waters's translation of this chapter from 2015 (Weber 2015 [1921], 75–127). For other chapters from *Economy and Society*, I rely on the older translation by A. M. Henderson and Talcott Parsons (Weber 1964 [1947]).

⁴ In the German original, Weber uses the term *Herrschaft*, which is most commonly translated as 'domination.' Strictly speaking, however, this translation is inaccurate because 'domination' carries a connotation of oppression while the German *Herrschaft* does not. Furthermore, it is worth pointing out that the aforementioned chapter is not the only place where Weber discusses the nature of bureaucracies. For a more comprehensive overview, see Höpfl (2006: 9–13).

bureaucracy. Thus, every office and, consequently, every officer is responsible for one type of task that is meant to contribute to accomplishing the shared task of the bureaucracy at large. Furthermore, Weber's point about the division of labor entails that there is only a very limited amount of 'overlapping' responsibilities among the individual officers.

- 2) Weber claims that the offices within a bureaucracy are organized according to "the principles of 'hierarchy of offices' and 'proper channels'" (Weber 2015 [1921], 77). Thus, the offices within a bureaucratic social system stand in a hierarchical relationship insofar as the officers within higher-ranked offices give *orders* to the officers in the lower-ranked offices. Paradigmatically, officers from lower-ranked offices only accept orders that have been passed through these 'proper channels.' While Weber does not understand this hierarchy primarily in epistemic terms, many bureaucracies do organize the flow of epistemic goods in a hierarchical way too. For example, in many bureaucracies, subordinate officers pass on a summary of relevant epistemic goods to their superiors but not to any officers of equal or lower rank.
- 3) Weber claims that officers within a bureaucracy must keep *written* records of their work (Weber 2015 [1921], 77).⁵ This principle is meant to ensure that *any* competent officer with the relevant permissions can take over the office at a later time. Within a proper bureaucracy, no officer depends on epistemic goods that are transmitted through informal channels. (This feature will turn out to be especially important for an epistemic assessment of bureaucracies.)
- 4) Weber claims that "the work as a *Beamte* [officer] typically requires an in-depth specialist training to undertake all specialized tasks of the *Amt* [office]" (Weber 2015 [1921], 78). In other words, every office within a bureaucracy is occupied only by officers who have received specialized training that guarantees that they can handle the tasks that are assigned to their office.
- 5) Weber claims that "a full-fledged *Amt* [office] occupies all the professional energy of the *Beamte* [officer] to process its tasks, regardless of the fact that his mandatory working hours in the office can be strictly limited" (Weber 2015 [1921], 78). In other words, officers within a bureaucracy are conceived as *full-time employees*. As Weber later notes, this arrangement is meant to ensure that officers carry out their assigned tasks in strict accordance with internal rules and do not adhere to rules imposed by other organizations on which they are financially dependent (Weber 2015 [1921]: 96).⁶
- 6) Weber claims that "the duties of the *Amt* [office] undertaken by the *Beamte* [officer] are based on general learnable rules and regulations, which are more or less firm and more or less comprehensive" (Weber 2015 [1921], 78). In other words, a bureaucracy provides rules that prescribe how every office has to handle the task that has been assigned to them. Again, this rule

⁵ Presumably, in Weber's time, *Schriftstücke* ('written records') simply stood in contrast to oral records. Thus, there is no reason to assume that 'written records' are meant to exclude, say, digital records.

⁶ This strikes me as a less compelling aspect of Weber's account. There are countless bureaucratic organizations that rely extensively on part-time employees. For example, many universities and colleges in the US employ part-time instructors (adjunct professors, PhD students, etc.) to handle a significant share of their teaching responsibilities. Regardless, the prevalence of part-time employment does not intuitively call into question their status as 'bureaucracies.'

is supposed to make sure that the offices within a bureaucracy operate in the same way independently of the individual officers that occupy them.

Weber's account of bureaucracies is not uncontroversial. Höpfl (2006, 9–18) laments that Weber provides differing accounts of bureaucracies on different occasions and often marks different features as 'essential.' He criticizes "the absence [...] of connective tissue linking together the various 'components' of bureaucracy, and the lack of prioritisation among them, despite the proliferation of headings, numbered paragraphs, sub-paragraphs and *Sperrschrift* (equivalent to italics)" (Höpfl 2006, 11). The inconsistencies are evident, for example, in the way Weber characterizes officers within a bureaucracy. Sometimes, he claims that it is an "essential" (Weber 1947 [1921], 335) feature of bureaucracies that officers offer their services as a means to receive a wage. On other occasions, he denies this fact outright and claims that officers within a bureaucracy must be driven by a sense of duty instead (Weber 2015 [1921], 79). To give another example, Weber frequently mentions secrecy as a crucial feature of bureaucracies: "the bureaucracy hides its knowledge and its conduct from criticism as well as it can" (Weber 2015 [1921], 116). Yet, it remains unclear how secrecy can be derived from the aforementioned list of six characteristic paradigmatic features.

I will not attempt to develop an account of bureaucracies that fixes these shortcomings of Weber's account. While this would be an interesting project for future research, Weber's account seems sufficient to provide a general idea of bureaucratic social systems. Since my goal here is to assess bureaucracies from an *epistemic perspective*, I will instead focus on enumerating a number of features that bureaucracies paradigmatically exhibit and that are relevant for analyzing their epistemic functioning:

- (a) **Shared Task Feature:** Bureaucracies pursue a shared task.
- (b) **Division of Labor Feature:** Bureaucracies are highly structured social systems insofar as their members have been assigned a social role that require them to carry out a task
 - (i) that contributes to accomplishing the shared task of the bureaucracy at large,
 - (ii) that is not identical to the task assigned to any other member of the bureaucracy, and
 - (iii) that is supposed to be carried out by following organization-wide rules.
- (c) **Epistemic Hierarchy Feature:** Bureaucracies organize the internal flow of epistemic goods in a hierarchical way.

I would like to emphasize that this is not meant to be a *definition* of bureaucracies. This is meant to be an incomplete list of epistemically relevant features that bureaucracies typically exhibit. I assume that this characterization is compatible with a range of different Weber-inspired definitions of bureaucracy.

Furthermore, this characterization of bureaucracies is meant to remain neutral with respect to some of the more controversial *political* aspects of Weber's account. As Mayntz (1965) notes, it is often overlooked that Weber understands bureaucracy as a particular form of *domination* [Herrschaft]. For my purposes, it is not necessary to burden the concept of bureaucracy with this political assumption. Accordingly, I do not need to follow Weber's distinction between 'officers' and 'non-officers' as well. For Weber, an officer is an employee who exerts authority in virtue of their position within the organization. Weber claims that neither ordinary rank-and-file employees nor low-level office workers (e.g., secretaries, typists, etc.) count as officers. Beetham (1997, 13) plausibly objects that this exclusion is hard to justify because both

types of workers typically *do* exert authority in virtue of their position within the bureaucracy. To avoid this controversy, I simply use the term ‘members’ for the employees of a bureaucracy (regardless of whether they count as ‘officers’ in Weber’s sense or not).

Let me explain the three aforementioned epistemically relevant features of bureaucracies in some more detail:

Shared Task Feature: A standard way to explain what it means for a social system to engage in a shared task is to claim that it must exhibit the following four features: first, the members of the social system have a ‘shared intention’ to achieve a certain goal; second, a number of ‘sub-plans’ have been assigned to every member of the social system in order to achieve this overarching goal; third, the actions carried out by the individual members of the social system are coordinated in an interactive way to achieve this overarching goal; and fourth, every member of the social system must make at least *some* contribution toward the overarching goal of the social system.⁷ This four-fold characterization of shared agency mostly seems correct to me. However, the first requirement seems hard to justify if we take into account that some shared tasks are carried out by social systems with *alienated* members, as Shapiro (2014) has pointed out. After all, we would say that a social system engages in a ‘shared task’ even if some of its members do not have the intention to achieve it. For instance, alienated workers in a company might lack the intention to achieve the goal of generating profit for the shareholders, but we would still say that the company constitutes a social system that pursues precisely this shared task. This objection is especially relevant for the discussion of bureaucracies because many organizational sociologists have argued that this type of alienation frequently occurs in bureaucracies (Tullock 1965; Downs 1966, 118–121).

Division of Labor Feature: I would like to clarify that the Division of Labor Feature is compatible with overlapping responsibilities to some extent. In certain cases, a bureaucracy may assign the same task to multiple members as a form of quality control. However, the Division of Labor Feature requires that the overall set of tasks assigned to different members of a paradigmatic bureaucracy are not identical. This strikes me as a modest assumption because assigning *identical* tasks to multiple members would indeed be highly inefficient and undermine the very idea of the division of labor. Furthermore, it is worth highlighting that the aforementioned Division of Labor Feature requires that the members of a bureaucracy paradigmatically cannot choose by themselves how they want to carry out their assigned tasks, but that they must do so by following official rules of the organization.

Epistemic Hierarchy Feature: By ‘hierarchical flow of epistemic goods,’ I mean that members of a bureaucracy occupy different ranks and that their assigned tasks correspond to these ranks. Lower-ranked members are assigned tasks that include sharing certain epistemic goods with relevant higher-ranked members, while higher-ranked members are assigned tasks that involve processing the epistemic goods passed on by their subordinates. The Epistemic Hierarchy Feature does not specify which kinds of epistemic goods must be organized hierarchically within a bureaucracy. These epistemic goods may include reports about members’ work, measures of their performance, or simply the outcomes of their assigned tasks. Thus, this assumption is compatible with Weber’s more specific claim that all members of a bureaucracy must keep written records of their activities.

⁷ For further discussion of shared agency, see Gilbert (1989), Bratman (1992), and Tuomela (2007).

I believe that this account of bureaucracies can gracefully handle two borderline cases:

First, there are some bureaucracy-like social systems whose shared task has no epistemic component. For example, consider Smith's pin factory again. The shared task of a pin factory is to produce pins and, in principle, one can produce pins without producing epistemic goods. I assume that these social systems do not count as paradigmatic cases of bureaucracies because their shared task has no distinctive epistemic component. A social system can only satisfy the Epistemic Hierarchy Feature if its shared task has at least *some* epistemic component. So, if a pin factory is to count as a genuine bureaucracy, it must at least demand that its members log or measure their activities and share the resulting epistemic goods with higher-ranked members within the organization.

Second, there are some social systems whose shared task has an epistemic component but that lack an internal hierarchy for epistemic goods. For instance, consider a beehive. The bees within a hive collectively pursue certain tasks (e.g., locating food sources and raising offspring) and accomplish these tasks by dividing labor (e.g., with queens, workers, and drones performing distinct social roles). Furthermore, bees even share information about the location of food sources with one another by performing a so-called 'waggle dance.' The direction and duration of this dance encodes the location and distance of the food source relative to the sun. For this reason, a beehive exhibits many of the aforementioned paradigmatic features of a bureaucracy. Yet, it does not intuitively qualify as a bureaucracy. My account can explain why: a beehive does not satisfy the Epistemic Hierarchy Feature because bees do not share epistemic goods in a hierarchical way. The waggle dance serves to make epistemic goods democratically accessible to other bees in the hive. There are no bees that have privileged access to these epistemic goods. For this reason, my account can accommodate the intuition that a beehive falls short of a paradigmatic case of a bureaucracy.

Finally, I would like to point out that my account of bureaucracies is applicable to both cases depicted at the beginning of this article. Both de Prony's team in *SUCCESS STORY* and the military bureaucracy in *FIASCO* exhibit all three features that my account ascribes to paradigmatic instances of bureaucracies:

First, de Prony's team (a) pursues a shared task, i.e., the task of computing accurate decimal approximations of logarithmic and trigonometric functions. Furthermore, his employees (b) try to accomplish this shared task by breaking it into smaller ones and distributing them in the team. In de Prony's team, some workers were assigned the role of 'planners.' They "took the basic equations for the trigonometric functions and reduced them to the fundamental operations of addition and subtraction" (Grier 2006, 36). Others were assigned the role of 'computers.' These computers carry out the fundamental operations of addition and subtraction, as requested by the planners. The planners then check the results provided by the computers by calculating the differences between adjacent values (Grier 2006, 37). Both computers and planners are supposed to carry out their respective tasks by following certain official rules. This description makes evident that de Prony's team exhibits (c) a clear epistemic hierarchy as well because the computers must send their results to the responsible planners.

Second, the military bureaucracy in *FIASCO* (a) pursues the shared task of defending US citizens against threats (which involves the epistemic task of monitoring threats to US citizens). Its members (b) try to accomplish this task by dividing it into smaller ones and distributing them among the members of the

organization. For instance, senior officers are responsible for broad strategic assessments, while rank-and-file soldiers focus on monitoring radar signals. Furthermore, there are official rules that determine how these tasks are meant to be carried out. There is (c) a clear epistemic hierarchy as well because the rank-and-file soldiers are tasked to report suspicious radar signals to their assigned superior officers.

§ 3. Methodological Considerations for the Epistemic Appraisal of Bureaucracies

The method of this investigation can be described as what Goldman calls ‘systems-oriented social epistemology’ (SYSOR-SE). As he puts it, ‘systems-oriented social epistemology’ studies social systems “in terms of their effects on epistemic outcomes” (Goldman 2011, 18). This method differs both from the ‘individual doxastic agent social epistemology’ (IDA-SE) and the ‘collective doxastic agent social epistemology’ (CDA-SE) that have dominated the field of social epistemology so far. In contrast to individual doxastic agent social epistemology, systems-oriented social epistemology is not necessarily concerned with simple social systems that consist only of a dyad of speaker and hearer. In contrast to collective doxastic agent social epistemology, systems-oriented social epistemology is not necessarily concerned with collective doxastic states such as collective beliefs or collective knowledge. So, in this article, I assess the expectable epistemic outcomes of a range of social systems that are more complex than simple speaker-hearer dyads and I will consider *all* relevant expectable epistemic outcomes (and not just the collective doxastic states that they exhibit).⁸ This includes epistemic values that are instantiated on the level of the *members* as individual doxastic agents and on the level of the *social system at large* as a collective doxastic agent.

When we engage in systems-oriented social epistemology, we have to make certain assumptions about epistemic normativity. I assume that the epistemic goodness of a social system depends on how it realizes different epistemic values. The most prominent candidates for such epistemic values are truth, knowledge, and understanding:

- (a) **Truth:** Roughly, when we are in possession of ‘truth,’ we represent an object as what it really is. Everyone should agree to that much, even if there is much disagreement about the nature of truth (Blackburn 2018).
- (b) **Knowledge:** There is widespread consensus that a justified true belief counts as knowledge if some further conditions are met, but it is notoriously controversial what these further conditions are. Almost everyone agrees that knowledge must satisfy some kind of ‘anti-luck’ condition in order to rule out Gettier-style counterexamples (Gettier 1963). Furthermore, although this fact is often explained in different ways, it is widely held that knowledge is *ceteris paribus* epistemically better than merely true belief (e.g., it is better to know the road to Larissa than to make a lucky guess that happens to be true) (Zagzebski 2003; Kvanvig 2003; Greco 2010, ch. 6).

⁸ Goldman (2011) leaves the relations between IDA-SE, CDA-SE, and SYSOR-SE somewhat unclear. I understand IDA-SE as a *type* of SYSOR-SE because IDA-SE studies a certain *type* of social system: namely, simple dyadic social systems. Similarly, I regard CDA-SE as a *type* of SYSOR-SE because it deals with a certain *type* of social system: namely, highly structured social systems that exhibit their own collective doxastic states. For this reason, the results of IDA-SE and CDA-SE are relevant to SYSOR-SE as well.

- (c) **Understanding:** Roughly, having understanding means “being able to ‘grasp’ or ‘see’ how the various parts of the world were systematically related” (Grimm 2012, 103). Furthermore, it is widely accepted that an agent who understands a certain subject matter should be able to provide *explanations* about that subject matter. Again, there is disagreement about the nature of understanding though. While some philosophers claim that understanding is an epistemic value that cannot be reduced to knowledge (Kvanvig 2003, 196–200; Elgin 2017, ch. 3), others regard it as a certain type of knowledge. For example, Greco’s account belongs to the latter category. Drawing inspiration from Aristotle, he argues that “understanding consists in a *systematic* knowledge of dependence relations” (Greco 2020, 133).

There is another methodological complication about epistemic normativity that we have to keep in mind in this context. Goldman (1988) points out that the epistemic value of truth alone can be instantiated in different ways (namely, ‘reliability,’ ‘power,’ and ‘speed’) that are all relevant for the epistemic evaluation of a social system. Arguably, this insight can be generalized to other epistemic values as well. I will briefly explain Goldman’s distinction between reliability, power, and speed, and elaborate on their counterparts for other epistemic values such as knowledge and understanding. First, we consider a social system *reliable* when it exhibits “a tendency to produce a high truth ratio of beliefs” (Goldman 1988, 26). More specifically, we can define reliability as the ratio of true beliefs to beliefs overall that a social system tends to produce. Thus, the more reliable a social system is, the less likely will it give rise to any false beliefs. Second, Goldman defines *power* as follows: “Power is the capacity of a process, method, system, or what have you to produce a large number of true beliefs” (Goldman 1988, 27). Even though power and reliability are both linked to truth, they are distinct standards of epistemic evaluation. On the one hand, a social system can be very reliable insofar as it tends to produce almost exclusively true beliefs, but at the same time lacks power insofar as it only produces few true beliefs overall. On the other hand, a social system can be very powerful insofar as it produces a large number of true beliefs, but lack reliability insofar as it also yields a significant number of false beliefs at the same time. The third dimension of truth-related epistemic goodness is *speed*, i.e., “speed in getting true beliefs” (Goldman 1988, 27). In general, we can define speed as the ratio of true beliefs that a social system tends to produce in a given time unit. We must distinguish speed from reliability and power. For example, a social system can be reliable but slow insofar as it tends to produce nothing but true beliefs but needs a long time to obtain each true belief. Arguably, Goldman’s threefold distinction can be generalized to other epistemic values as well. If X is an epistemic value, we can typically distinguish three dimensions of epistemic goodness that are analogous to reliability, power, and speed. First, it is usually *prima facie* epistemically good to be disposed to instantiate X with a low rate of failures (which corresponds to ‘reliability’). Second, it is usually *prima facie* epistemically good to be disposed to produce many instances of X (which corresponds to ‘power’). Third, it is usually *prima facie* epistemically good to be disposed to produce many instances of X in a given time unit (which corresponds to ‘speed’).

We saw that bureaucracies are highly structured social systems (which follows from the Division of Labor Feature mentioned in section § 2). Thus, every member of a bureaucracy has a distinct social role. As we have seen, the social role assigned to every member has at least some epistemic component, which in turn impacts the beliefs that the individual members can be expected to hold. In particular, we can assume that the members of a bureaucratic social system are prone to form two types of beliefs:

- 1) **Organizational Beliefs:** First, the Division of Labor Feature requires that the members of a bureaucracy have beliefs about their own position and tasks within the organization. They are rarely expected to understand the internal structure of the entire social system, but should know the bureaucracy-wide rules governing their assigned tasks. In particular, because of the Epistemic Hierarchy Feature, lower-ranked members must form beliefs that identify the appropriate superiors to whom epistemic goods should be reported and higher-ranked members must form beliefs that identify the subordinates from whom such goods should be received. For example, in *FIASCO*, any rank-and-file radar operator should have certain beliefs about, say, whom they must inform in case they encounter a suspicious radar signal. Henceforth, I will call these kinds of beliefs ‘organizational beliefs.’

- 2) **Content Beliefs:** Second, because of the Epistemic Hierarchy Feature, the members of any bureaucratic social system will form certain beliefs in virtue of the epistemic tasks that they have been assigned. Keep in mind that we saw in section § 2 that even the members of bureaucracies that do not directly pursue an epistemic task (e.g., workers in a pin factory) must engage at least in *some* epistemic tasks such as logging or measuring their activities. Otherwise, these social systems cannot be considered paradigmatic instances of bureaucracies. Henceforth, I will call beliefs that emerge in this way ‘content beliefs.’ These content beliefs stem from two sources. On the one hand, members will form certain beliefs *directly* by carrying out the epistemic tasks that have been assigned to them. For instance, in *SUCCESS STORY*, a ‘computer’ might form the belief that $349,827 + 481,265 = 831,092$; and a rank-and-file radar operator in *FIASCO* will form the belief that there is a suspicious radar signal on December 7, 1941 at 7:02 am. On the other hand, members will form certain beliefs *indirectly* because it is part of the responsibility of other members to make the results of their assigned epistemic tasks accessible. For instance, a ‘computer’ in *SUCCESS STORY* will tell the responsible ‘planner’ that they calculated that $349,827 + 481,265 = 831,092$ and the ‘planner’ will form the belief that this is true.

Thus, every member of a bureaucracy can be expected to form certain role-specific organizational beliefs and content beliefs. Furthermore, we have seen that if appropriate conditions are satisfied, then these beliefs will count as ‘knowledge’ as well; and if further conditions are satisfied, then these beliefs will also amount to ‘understanding.’ Beyond the level of individual cognition, bureaucratic systems may also instantiate collective doxastic states (e.g., collective knowledge). So, when I provide a systems-oriented social epistemology of bureaucratic social systems in this article, I evaluate the epistemic goodness of the respective mix of true beliefs, knowledge, and understanding of individual members and the bureaucracy at large as a collective doxastic agent.

§ 4. Ideal Social Epistemology of Bureaucracies

Weber himself is keen to avoid evaluating bureaucracies *simpliciter* (following his commitment to ‘value-free sociology’). Nonetheless, he attributes several noteworthy epistemic advantages to these social systems. As I briefly noted at the very beginning of this article, Weber argues that the distinctive role of trained specialists within bureaucracies makes these social systems epistemically superior to social systems in which domination is based on tradition or charisma: “Bureaucratic administration means fundamentally the exercise of control on the basis of knowledge” (Weber 1964 [1947], 339). Furthermore, Weber argues that “[a] fully developed bureaucratic mechanism compares to other administrative forms in

the same way a machine compares to nonmechanical means for producing goods” (Weber 2015 [1921], 96). Thus, just as machines tend to produce goods more quickly and reliably than other forms of production, bureaucracies can be expected to outperform other social systems in their capacity to generate epistemic goods in a fast and reliable way.

Similarly, if we just stick to the list of epistemically relevant features of bureaucracies provided in section § 2, we might come to the conclusion that bureaucratic social systems are overall epistemically good for the following three reasons:

- 1) **Reliable and Fast Formation of Organizational Beliefs:** Because of the Division of Labor Feature, we can expect that members of bureaucracies have been assigned a determinate task and that there are clear rules that determine how they are supposed to carry out their assigned task. For these reasons, we may expect that members of bureaucracies will mostly form true organizational beliefs. Therefore, members of bureaucracies can be expected to exhibit a high level of *reliability* with respect to organizational beliefs. For the same reason, we can expect that members of bureaucracies will be able to form true organizational beliefs *quickly*. In contrast, in social systems that distribute labor by means of informal arrangements, individual members must navigate a complex network of social relationships to determine what tasks they are expected to perform and how they are supposed to perform them (e.g., consider how complicated it often is to figure out who is supposed to contribute what to a shared birthday gift in a friend group).
- 2) **Reliable and Fast Formation of Content Beliefs:** The Division of Labor Feature tells us that bureaucracies accomplish their shared task by dividing labor. It is well-known that the division of labor positively impacts *productivity* in general. Adam Smith has pointed out that there are three reasons for this. First, when labor is divided in a social system, every member is solely responsible for one kind of task. Therefore, we can expect that every member’s “dexterity” (Smith 1976 [1776], 18) in carrying out their assigned task will improve over time. Second, the division of labor facilitates “saving [...] time commonly lost in passing from one sort of work to another” (Smith 1976 [1776], 18). Third, since the division of labor makes every member solely responsible for certain tasks, we can expect that the members will become better at identifying and utilizing *technical aids* that can help them in carrying out these tasks (Smith 1976 [1776], 19–20). Because of the Epistemic Hierarchy Feature, we know that the tasks assigned to members of bureaucracies always have at least some epistemic component. Thus, Smith’s insights into the productivity-enhancing role of the division of labor seem applicable to epistemic labor as well: we may expect that the division of labor will improve members’ *reliability* with respect to content beliefs, i.e., the ratio of true content beliefs they tend to produce, and their *speed*, i.e., the ratio of true content beliefs that they tend to produce in a given time unit.
- 3) **Ability to Tackle Complex Epistemic Tasks:** Third, because of the Division of Labor Feature, bureaucracies allow us to tackle complex epistemic tasks that exceed the cognitive limitations of an isolated individual agent. Many bureaucratic organizations engage in research projects that require far more epistemic labor and specialized disciplinary training than a single researcher could accomplish in a lifetime. For example, consider the Human Genome Project. The Human Genome Project was an enormous research project that ran from 1990–2003 and aimed to

determine the DNA sequence of the entire human genome. The International Human Genome Sequencing Consortium (IHGSC) divided the genome into segments that were assigned to 20 universities and research centers worldwide. Each research team sequenced their assigned regions and contributed the results to a shared database (National Human Genome Research Institute 2024). This bureaucratic organization allowed researchers to form true beliefs about the structure of the human genome for the first time in history. For this reason, bureaucracies can positively impact *power* (in Goldman’s sense) with respect to content beliefs, i.e., the quantity of true content beliefs that are accessible to their members overall. While bureaucracies may not be the only type of social system with this epistemic advantage, it is undoubtedly one of them.

Thus, we can expect that bureaucracies will positively impact their members’ reliability and speed with respect to both organizational and content beliefs. Furthermore, bureaucracies potentially have a positive impact on their members’ power with respect to content beliefs as well. Moreover, since bureaucratic social systems make their members prone to form true beliefs in a way that is justified and non-coincidental, we can expect that they facilitate the formation of *knowledge* as well. Thus, on these grounds, one might come to the conclusion that bureaucracies are overall epistemically good social systems.⁹ For instance, *SUCCESS STORY* can be used to illustrate this verdict: because de Prony’s team has clear rules that determine the workers’ responsibilities, both ‘computers’ and ‘planners’ manage to form organizational beliefs reliably and quickly. Furthermore, de Prony’s team could actually compute decimal approximations of logarithmic and trigonometric functions with a previously unprecedented level of reliability and speed. Yet, as we will see in the next section, in many real-life scenarios (e.g., *FIASCO*), bureaucracies give rise to less favorable epistemic outcomes than this ideal social epistemology suggests.

§ 5. Non-Ideal Social Epistemology of Bureaucracies

Many critics have objected that Weber’s account of bureaucracies is oversimplified and claim that it presents bureaucratic social systems in an overly favorable light (Gouldner 1954; Merton 1968 [1940]; Simon & March 1993 [1958]; 36–47). In many cases, bureaucracies hardly resemble Weber’s image of a ‘machine’ in which all parts fit neatly together; rather, they are better described by Kafka’s image of a ‘castle,’ whose inner workings remain opaque and mysterious.

In this section, I offer a similar objection against the social epistemology of bureaucracies that we encountered in the previous section. I argue that the aforementioned theory of bureaucracies qualifies as an *ideal theory* in two related senses. First, it abstracts from many epistemically relevant features of real-world bureaucracies. Second, as a consequence, it is also ‘ideal’ in a normative sense, i.e., it portrays bureaucracies as epistemically better than is warranted. While a high level of abstraction is not *per se* a problem, it is unfortunate insofar as the aforementioned ideal theory is unable to explain why many real-life bureaucracies give rise to unfavorable epistemic outcomes. For instance, the aforementioned

⁹ Some prominent sociologists claim that another potential epistemic advantage of bureaucratic social systems is their ability to “keep valuable information from economic competitors or from foreign and potentially hostile political groups” (Merton 1968 [1940], 561; cf. Weber 2015 [1921], 116–118). Indeed, many bureaucracies have institutionalized internal constraints on the sharing of epistemic goods. For example, military bureaucracies often operate according to the ‘need to know’ principle, which restricts access to epistemic goods in order to prevent leaks to hostile doxastic agents. However, it is not clear that secrecy counts as a paradigmatic feature of bureaucratic social systems (at least, if we presuppose the account of bureaucracies from section § 2).

ideal social epistemology of bureaucracies struggles to explain why the military bureaucracy portrayed in *FIASCO* yields such poor epistemic outcomes.¹⁰

In this section, my goal is to develop a non-ideal systems-oriented social epistemology of bureaucracies. In other words, I aim to develop a more realistic account of the epistemic outcomes that these social systems can be expected to give rise to. I will show that the ideal social epistemology of bureaucracies depicted in the previous section relies on a superficial characterization of (a) their *human components*, (b) their *non-human components*, and (c) their *environment*.¹¹ For this reason, bureaucracies in real life can be expected to give rise to lower levels of reliability and speed with respect to both organizational and content beliefs than the ideal theory suggested. Furthermore, as I will show, another limitation of this ideal theory is (d) its overly narrow conception of relevant *epistemic goods*: namely, I show that bureaucracies tend to obstruct the formation of *understanding*, which is a serious epistemic downside that emerges even under favorable conditions. I do not claim that paying attention to these aspects will give us a complete picture of how bureaucracies operate under realistic conditions; yet, I hope that it is a step in the right direction that allows us to form a more realistic understanding of the epistemic functioning of these social systems.

Even if the ideal theory of bureaucracies presented in the previous section struggles to account for the epistemic outcomes of many real-world bureaucracies, it is far from insignificant. This ideal theory illustrates the epistemic potential of bureaucratic social systems, i.e., it shows what kinds of epistemic goods bureaucracies can produce under favorable conditions. In this way, the ideal theory serves an important orienting function: it allows us to make an informed decision whether it is worth creating or maintaining the conditions under which bureaucracies realize these epistemic advantages.¹²

§ 5.1. Human Components

In this section, I argue that the ideal systems-oriented social epistemology of bureaucracies depicted in section § 4 relies on a superficial characterization of the *human components* of these social systems. We saw that bureaucracies are highly structured social systems, i.e., they assign a distinct social role to every member. However, it is a persistent problem of real-life bureaucracies that even the most disciplined members of bureaucracies do not always adhere to their assigned social roles:

“In practice [...] people’s personalities are never so totally subsumed into their roles. They come to the organization as individuals, with personal needs and expectations for which they seek satisfaction: from social intercourse at the workplace; from the exercise of skill and a measure of control over the work process; from being treated ‘as people’ rather than as the impersonal occupants of a role.” (Beetham 1997, 16)

¹⁰ For a more elaborate discussion of the virtues and vices of ‘ideal theory,’ see Mills (2005), Stahl (2024), and Kremers (forthcoming). McKenna (2023) discusses the role of ideal theory in epistemology specifically.

¹¹ Similarly, McKenna distinguishes between four different conceptions of ideal theory in epistemology: “(1) Idealizations about agents and the interactions between them. (2) Idealizations about human psychology. (3) Idealizations about social institutions. (4) Idealizations about our environment” (McKenna 2023, 32). McKenna seems to overlook that an overly narrow focus on certain epistemic goods such truth can constitute an ideal epistemic theory as well.

¹² I have argued in Kremers (forthcoming) that it is an underappreciated practical aspect of ideal theories in general that they can help us to clarify which ideals are actually worth striving for.

These conflicts between assigned social roles and members' personal interests give us a reason to temper our optimism about the epistemic performance of bureaucracies for at least two reasons:

First, we saw that members of bureaucracies tend to form beliefs *indirectly* because it is part of the assigned social role of other members to obtain certain epistemic goods and make them accessible. However, if other members are prone to deviate from their assigned social role because of conflicting personal interests, then we may expect that they will also forge, omit, and distort the contents of these epistemic goods if doing suits their personal interests. For this reason, members of a bureaucracy will be more prone to form false beliefs based on the testimony of other members. As a result, these members' reliability and speed will fall short of the high levels that were promised by the ideal social epistemology presented in the previous section. For example, the phenomenon of the 'blue wall of silence' can be used to illustrate this point: it is a well-known problem that police officers in the US are reluctant to report misconduct committed by fellow officers because doing so conflicts with their personal interests (Chin & Wells 1998). Thus, because police officers tend to neglect certain aspects of their assigned social role, members of the police bureaucracy turn out to be more prone to form false beliefs about the quantity and quality of police misconduct.

Perhaps, one might object that deviations from the assigned social roles might actually *enhance* the overall epistemic performance of a bureaucracy. As a matter of fact, many organizational sociologists take informal social norms to be indispensable for the proper functioning of bureaucratic organizations. It is well-known that carrying out assigned tasks strictly 'by the book' while ignoring informal social arrangements often amounts to a form of sabotage (Gouldner 1954, 175–176; Luhmann 1972 [1964], 301). However, I would like to clarify that this objection does not call into question the point I made. I have argued that the ideal theory's optimism about the epistemic performance of bureaucracies is misplaced: we have no good reason to believe that bureaucracies must generally exhibit high levels of reliability and speed. However, this position does not commit me to the view that bureaucracies must generally exhibit low levels of reliability and speed either. There are indeed some bureaucracies in which informal social norms play a crucial role and these bureaucracies might indeed give rise to good epistemic outcomes as a result. However, in such cases, the epistemic success is due to these informal arrangements rather than to their bureaucratic structure as such.

Second, these conflicts between official rules and personal interests are responsible for another epistemic downside of bureaucracies. If such conflicts abound, individual members are likely to rely on informal social norms rather than formal ones. However, informal norms have the disadvantage that they require considerable time and effort to learn.¹³ For this reason, we have little reason to assume that bureaucracies make it especially easy for their members to form true beliefs about their own roles and tasks within the organization. Thus, the ideal theory of bureaucracies overstates its case when it claims that bureaucracies must exhibit a high level of reliability and speed with respect to organizational beliefs as well. Presumably, this will seem familiar to many because many bureaucracies are described by their members as chaotic and difficult to navigate. For example, this is a central theme of Kafka's novels on bureaucracy

¹³On the interaction between formal rules and informal norms, see Downs (1966, 51–68). Gouldner's *Patterns of Industrial Bureaucracy* (Gouldner 1956) provides a classic case study of this dynamic.

(Kafka 2009 [1926]; Hodson et al. 2012) (who, after all, was a long-standing employee of the Workers' Accident Insurance Institute for the Kingdom of Bohemia).

There are a number of risk factors that facilitate these conflicts between members' personal interests and their assigned social role within the bureaucracy:

- a) First, the *size* of a bureaucracy is a major risk factor. On the one hand, large social systems struggle to detect violations of social norms, which makes compliance with them generally less attractive to its members. Thus, in large bureaucracies, it becomes easier to 'slip through the cracks' and get away with ignoring one's social role (Olson 1965, 60–65). This applies to the epistemic aspects of one's assigned social role as well. For this reason, we can expect that members of large bureaucracies will be more prone to omit, distort or falsify sensitive epistemic goods and that the epistemic performance of the organization at large will suffer as a result.

There is another related reason why the size of a bureaucracy negatively impacts its epistemic performance: namely, large bureaucracies tend to exhibit steeper hierarchies. Therefore, epistemic goods must pass through more levels before reaching those at the top. At each stage, these epistemic goods are subject to filtering by members who may have incentives to distort or falsify them if doing so suits their personal interests. As a result, the epistemic goods available to those members at the top of a large bureaucracy are the product of multiple such filtration processes, which ends up degrading their quality (Williamson 1967, 126–127; Simon 1998 [1947], 215).

- b) Second, the presence of *administrative burdens* is another risk factor that facilitates conflicts between members' personal interests and their assigned social role. The process of filing and retrieving reports in a bureaucracy requires much time and effort. These administrative burdens can become so overwhelming that individual members will not fulfill their social roles as intended.¹⁴

Electronic Health Records (EHRs) nicely illustrate this risk. In many countries, EHR systems have been adopted to replace paper-based medical records with the aim of improving the accessibility of patient information and supporting clinical decision-making (e.g., by providing drug interaction alerts). In practice, however, EHRs are associated with substantial administrative burdens: they require physicians to enter large amounts of patient information in cumbersome, structured formats and to navigate complex user interfaces in order to retrieve such information. In one study, physicians spent nearly twice as much time on EHR-related tasks as on direct patient interaction (Arndt et al. 2017). There is evidence that these administrative burdens lead physicians to deviate from fulfilling their assigned social roles as intended: on the one hand, EHRs make it difficult for them to retrieve clinically relevant information because electronic

¹⁴ For further discussion of 'administrative burdens,' see Herd and Moynihan (2018). Incidentally, it has been recognized that this feature of bureaucracies can be exploited to *sabotage* them as well (Gouldner 1954, 175–176; Luhmann 1972 [1964], 301). For example, the precursor of the CIA, the US Office of Strategic Services, gave its undercover agents the following advice in its now-declassified 'Simple Sabotage Field Manual' from 1944: "Insist on doing everything through 'channels.' Never permit short-cuts to be taken in order to expedite decisions. [...] When possible, refer all matters to committees, for 'further study and consideration.' Attempt to make the committees as large as possible—never less than five." (Office of Strategic Services 1944, 28)

records often contain large amounts of redundant or irrelevant data. On the other hand, extensive documentation requirements facilitate low-quality data entry (e.g., through the indiscriminate use of the copy-and-paste function) (Nijor et al. 2022). For these reasons, the high administrative burdens associated with EHRs undermine the medical bureaucracy’s ability to produce true beliefs about patients’ conditions.

- c) Third, the presence of *hostile doxastic agents* is another risk factor. When bureaucracies are surrounded by hostile doxastic agents, it is likely that these agents will attempt to disrupt the epistemic functioning of the social system by incentivizing members to deviate from their assigned social roles. This can occur through various means, such as corruption, coercion, or the spread of mis- and disinformation.

For example, Kolodny (2020) points out that many pharmaceutical companies acted as hostile doxastic agents over the past three decades by attempting to influence the US Food and Drug Administration’s (FDA) regulatory policies concerning opioid use. On the one hand, they accomplished this by spreading misinformation about the risks of opioid addiction (e.g., through studies with questionable research designs such as ‘enriched enrollment randomized withdrawal trials’). At the same time, the “revolving door between the FDA and the pharmaceutical industry” (Kolodny 2020) created incentives for individual FDA regulators to deviate from their official roles in order to increase their chances of securing lucrative private-sector positions after leaving public service.

§ 5.2. Non-Human Components

In this section, I argue that the ideal social epistemology depicted in section § 4 relies on a superficial characterization of the *non-human components* of these social systems as well. Most notably, I show that it tends to overlook the role of available *communication technology*.

We saw that bureaucracies only promised to facilitate the reliable and fast formation of content beliefs because producing true content beliefs from scratch is relatively costly. Thus, these epistemic benefits can manifest only if the epistemic labor required to confirm the truth of a belief by oneself exceeds the epistemic labor that is required to gather this information from a fellow member. However, this condition is not always met because the sharing of epistemic goods requires epistemic labor itself:

“Communication requires definite costs. Every message involves the expenditure of time to decide what to send, time to compose the message, the resource-cost of transmitting the message (which may consist of time, money, or both), and time spent in receiving the message. Also, if the message passes over a channel operating near its capacity, it may cancel or delay other messages.” (Downs 1966, 69)

The sharing of epistemic goods within bureaucracies does not happen in the void but is necessarily mediated by communication technology. Depending on the available communication technology, more or less epistemic labor will be required to retrieve an epistemic good that was previously produced by another member. Thus, depending on the quality of the available communication technology, reliability and speed with respect to content beliefs will be compromised. To give a simple example, if members of a bureaucracy can share information only verbally while working in a noisy office, then they will be prone

to form false testimonial beliefs (unless effective countermeasures are implemented).¹⁵ To give another example, an enormous amount of epistemic labor is required to extract epistemic goods encoded in jargon-laden legal terminology that is common in many bureaucracies. In both cases, the quality of the communication technology impacts the reliability and speed with which epistemic goods can be shared within a bureaucracy.¹⁶

Relatedly, the inherent limitations of different communication technologies force another epistemic problem on bureaucracies: due to the varying effectiveness of different communication technologies for different messages, we are forced to choose between optimizing for speed and optimizing for reliability. For instance, we can organize internal communication in terms of *forms*, i.e., standardized documents that specify exactly what information is required to perform a certain action. As a result, the transmission of epistemic goods within the social system will be reliable but slow. Yet, if we simplify these forms, we can improve the speed of transmission, but only at the expense of reliability because many unusual epistemic goods cannot be transmitted easily by means of simple forms (e.g., consider how unusual patient symptoms are misrepresented when reduced to a limited set of checkboxes).¹⁷ Nguyen even argues that the difficulty to represent unusual epistemic goods in terms of forms constitutes a form of *epistemic oppression* because “[t]hose whose ideas don’t fit comfortably into the regularized institutional language are at a significant disadvantage in participating in the production and dissemination of knowledge” (Nguyen 2021, 249).

§ 5.3. Environment

In this section, I show that the ideal systems-oriented social epistemology of bureaucracies depicted in section § 4 relies on a superficial characterization of the *environment* of these social systems as well.

We saw that the ideal theory promised that bureaucracies facilitate the reliable and fast formation of content beliefs because the division of labor allows the individual members to become more ‘dexterous’ in carrying out their assigned epistemic tasks and better at picking suitable technological aids. However, it is important to keep in mind that such a training effect can only emerge over time. The formation of true content beliefs within bureaucracies cannot be expected to be especially reliable and fast right from the start. The members must first have a chance to become familiar with their assigned epistemic tasks. This problem at least contributed to the poor epistemic outcomes of the military bureaucracy in *FIASCO*: on the day of the attack, Lieutenant Kermit Tyler supervised the information center at Fort Shafter for the very first time. Thus, he had virtually no chance to familiarize himself with his epistemic task and

¹⁵ On the role of noise in communication systems, see Shannon and Weaver (1964 [1949], 65 ff.). As Dennett (2017, ch. 9) notes, *digitization* is an effective measure to improve the reliability of information transmission under noisy conditions. For example, binary signals (e.g., 0 and 1) are easier to discriminate in noisy environments, but have the downside that more signals must be sent to transmit the same amount of information.

¹⁶ Incidentally, the discipline of economics went through a similar learning process. Adam Smith’s discovery of the productivity-enhancing potential of the division of labor in the 18th century made it tempting to believe that organizing *all* commodity production within a single giant firm would maximize productivity. However, economists in the 20th century realized that large firms are not always more productive because, as Ronald Coase puts it in an influential article, “as a firm gets larger, [...] the costs of organising additional transactions within the firm may rise” (Coase 1937, 394). Thus, just as *transaction costs* limit the productivity of commodity production within firms, *transmission costs* limit the productivity of knowledge production in bureaucracies.

¹⁷ Reverman (2025) analyzes this problem in more detail.

consequently misinterpreted the radar signal of the incoming attack. In *SUCCESS STORY*, by contrast, de Prony’s team faced a perfectly stable environment, namely that of the timeless realm of numbers.¹⁸

This is an important limitation of the ideal social epistemology of bureaucracies presented in section § 4 because it means that the epistemic goodness of these social systems is compromised in rapidly changing environments. In rapidly changing environments, individual members of bureaucracies will frequently be assigned different epistemic tasks. As a result, the training effects that promise to facilitate the formation of true content beliefs in bureaucratic social systems will be less pronounced. This conclusion aligns with a number of prominent findings on organizational sociology: for example, Burns and Stalker have argued that the bureaucratic form of organization delivers the most promising outcomes in stable environments (Burns & Stalker 1961, 119 ff.). Similarly, Merton has argued that the constituent rules of bureaucratic social systems often turn out to be cumbersome “under special conditions not clearly envisaged by those who drew up the general rules” (Merton 1968 [1940], 254).

§ 5.4. Epistemic Goods

Even if we keep the aforementioned limitations concerning human, non-human, and environmental factors in mind, the poor epistemic performance of the military bureaucracy in *FIASCO* remains hard to explain. To make sense of *FIASCO*, we must recognize that the ideal social epistemology of bureaucracies depicted in section § 4 relies on a superficial characterization of *epistemic goods* as well. Perhaps, bureaucracies fare well if we measure their epistemic goodness solely in terms of truth or knowledge; yet, as I will argue in this section, they tend to obstruct the formation of *understanding*.

We saw that bureaucracies pursue a shared task by dividing the labor necessary to accomplish it. As a result, the role-specific content beliefs that members of a bureaucracy can be expected to form will be *heavily compartmentalized*. In other words, members of a bureaucracy will likely have true beliefs about matters that are essential for fulfilling their assigned social role, but not about much else. Thus, they can be expected to gain knowledge about this isolated domain, but will struggle to obtain knowledge that would explain how this domain is embedded in a larger context. This has an important consequence for the epistemic value of *understanding*. If we define understanding as the “*systematic knowledge of dependence relations*” (Greco 2020, 133), then we see that *bureaucracies obstruct the formation of understanding*. For instance, in *SUCCESS STORY*, the individual ‘computers’ form many true beliefs about sums and differences. Yet, since their job consists of nothing but computing sums, they are hardly able to explain what exactly a logarithm or a trigonometric function is. As Grier points out, most of de Prony’s ‘computers’ had no noteworthy understanding of mathematics:

“Many of de Prony’s computers were former servants or wig dressers, who had lost their jobs when the Revolution rendered the elegant styles of Louis XVI unfashionable or even treasonous. They were not trained in mathematics and held no special interest in science. De Prony reported that most of them ‘had no knowledge of arithmetic beyond the two first rules [of addition and subtraction].’ They were little different from manual workers and could not discern whether they were computing trigonometric functions, logarithms, or the orbit of Halley’s comet. One labor historian has described them as intellectual machines, ‘grasping and releasing a single piece of ‘data’ over and over again.” (Grier 2005, 36)

¹⁸ I am grateful to B. Huebner for pressing me on this point.

As a bureaucracy grows larger and more complex, its members' content beliefs become increasingly fragmented, and their understanding can be expected to suffer accordingly (unless effective countermeasures are implemented).

That understanding tends to disintegrate in bureaucratic social systems is not necessarily a novel insight. As I already mentioned, Weber himself warns that bureaucracies can turn us into “specialists without mind [*Fachmenschen ohne Geist*]” (Weber 1993 [1930], 182, translation modified). Similarly, with his comment that the rise of bureaucracies resembles a “polar night of icy darkness and hardness, no matter which group wins the outward victory now” (Weber 1994 [1919], 368), he also indicates that he had mixed feelings about these social systems. Much earlier, Adam Smith pointed out that the division of labor in general leads to cognitive deterioration:

“The man whose whole life is spent in performing a few simple operations, of which the effects are perhaps always the same, or very nearly the same, has no occasion to exert his understanding or to exercise his invention in finding out expedients for removing difficulties which never occur. He naturally loses, therefore, the habit of such exertion, and generally becomes as stupid and ignorant as it is possible for a human creature to become. The torpor of his mind renders him not only incapable of relishing or bearing a part in any rational conversation, but of conceiving any generous, noble, or tender sentiment, and consequently of forming any just judgment concerning many even of the ordinary duties of private life.” (Smith 1976 [1776], 782)

In a similar way, Smith argues that members of ‘barbarous societies’ regularly find themselves in a wide range of situations, which allows them to develop a rich network of mental connections. This cannot be said of a worker whose life primarily consists of straightening wire in a pin factory. Therefore, Smith concludes that living in a society that does *not* rely on the extensive division of labor has important epistemic benefits (Smith 1976 [1776], 783).

The internal epistemic hierarchy of bureaucratic social systems can even exacerbate this effect. We saw that epistemic goods within bureaucracies flow hierarchically, i.e., members of lower rank are expected to share epistemic goods with their superiors. We already saw in section § 5.1 that epistemic problems emerge because individual members on every level of the hierarchy sometimes omit, distort, and falsify epistemic goods for personal reasons. Now, we can see that the epistemic hierarchy is responsible for a serious epistemic downside of bureaucracies even if none of its members have any personal interests that conflict with their assigned social roles. The underlying reason is that members of a bureaucracy must simply filter out some epistemic goods on every level of the hierarchy in order to avoid cognitive overload (Williamson 1967, 126–127). As a result, members in higher ranked positions *ceteris paribus* have access to fewer epistemic goods. This makes it difficult for them to form a systematic knowledge of dependence relations as well. Therefore, we can expect that their understanding will likely be impoverished as well.

I began this paper by pointing out that bureaucracies can give rise to favorable epistemic outcomes in cases like *SUCCESS STORY*, but also to unfavorable epistemic outcomes in cases like *FIASCO*. I believe that my non-ideal systems-oriented social epistemology of bureaucracies is finally sufficiently elaborate to explain the contrast between *SUCCESS STORY* and *FIASCO*. The epistemic shortcomings in *FIASCO* cannot be explained with limitations of human discipline or of communication technology alone. On the

one hand, the poor epistemic performance of the military bureaucracy in *FIASCO* was facilitated by the fact that the surprise attack occurred during Kermit Tyler's very first shift at the information center at Fort Shafter. As a result, he had only a few opportunities to train his ability to interpret radar signals at this location. Yet, most importantly, I believe that we can observe suboptimal epistemic outcomes in *FIASCO* because *understanding has disintegrated*: Kermit Tyler had knowledge of the suspicious radar blip but lacked knowledge of the current geopolitical situation that would allow him to recognize that the radar signal is caused by approaching enemy aircraft. In contrast, Tyler's commanding officers did have this knowledge, but were unaware of the radar signal. Thus, neither Tyler nor his commanding officers had knowledge of *both* the present radar signal *and* the broader intelligence necessary to understand its significance. For this reason, their understanding of the situation was severely compromised, which resulted in them failing to undertake any steps in response to the imminent attack. Of course, understanding has disintegrated in *SUCCESS STORY* as well. We saw that most of de Prony's workers do not even know what logarithmic and trigonometric functions are. Still, the epistemic value of understanding is only of minor importance for evaluating the epistemic success of de Prony's research bureaucracy. Its epistemic success is largely determined by its ability to produce true beliefs and knowledge about the decimal approximations of logarithmic and trigonometric functions. The fact that de Prony's workers only had a superficial understanding of mathematics in general does not impact how we evaluate the epistemic performance of this social system at large. In contrast, for the military bureaucracy in *FIASCO*, it is of the utmost importance that decision-makers have a proper understanding of the military situation. Thus, I conclude that my non-ideal systems-oriented social epistemology of bureaucracies is able to explain the favorable epistemic outcomes in *SUCCESS STORY* as well as the unfavorable epistemic outcomes in *FIASCO*, which was my main goal for this section.

Finally, one clarification is in order. I have painted a picture of bureaucracies that is much less rosy than the picture from section § 4. Perhaps, one might concede that bureaucracies do not make their individual members epistemically better off but insist that the bureaucracy at large is epistemically good because it exhibits certain collective doxastic states that are epistemically good (e.g., collective knowledge or collective understanding). Yet, I would like to clarify that nothing I have said is at odds with this view. In some cases, a bureaucracy might be considered overall epistemically good because it embodies certain epistemically valuable collective doxastic states even though its members would not count as epistemically good in their own right. Still, even in such a social system, the fact that its individual members do not embody epistemic values in isolation is something that makes the social system *prima facie* epistemically worse. I have not denied that these epistemically bad features of bureaucracies cannot be outweighed by other epistemically good features that they exhibit on a collective level.¹⁹

§ 6. Conclusion

Let me summarize. I began this article by depicting two historical case studies of bureaucracies: essentially, *SUCCESS STORY* was a best case scenario while *FIASCO* was a worst case scenario. My goal was to explain why bureaucracies lead to epistemic success in cases like *SUCCESS STORY* while they

¹⁹ To echo an analogy made by Dennett (2017, ch. 8), a human being consists of trillions of cells that can hardly be considered intelligent in their own right. Yet, taken together they constitute an intelligent being that can instantiate epistemic values such as truth, knowledge, and understanding. Even in this case, the fact that the individual cells lack intelligence remains *prima facie* epistemically bad, but this fact is outweighed by the epistemic values that the organism exhibits at large.

lead to serious epistemic problems in cases like *FIASCO*. First, I developed an ideal social epistemology of bureaucracies that explained why one might be inclined to think that these social systems generally give rise to good epistemic outcomes. I considered three reasons. First, because there are clear rules that determine how members are supposed to carry out their assigned tasks within bureaucracies, we can expect them to form beliefs about their own position and tasks within the organization in a reliable and fast manner. Second, because bureaucracies pursue a shared task that is at least partially epistemic and that is accomplished by dividing labor, we can expect that its members' reliability and speed with respect to these epistemic tasks will improve over time as well. Third, bureaucracies can tackle remarkably complex epistemic tasks that individual inquirers would be unable to take on, i.e., its members' power (in Goldman's sense) may improve as well. We saw that this ideal systems-oriented social epistemology of bureaucracies nicely illustrated what kinds of epistemic goods bureaucracies can produce under favorable conditions (e.g., in *SUCCESS STORY*). However, we saw that it cannot explain the epistemic shortcomings that some bureaucracies give rise to (e.g., *FIASCO*). To remedy these shortcomings, I developed a non-ideal systems-oriented social epistemology of bureaucracies. I showed that the ideal account of bureaucracies relies on a superficial characterization of the human and non-human components of these social systems, their environment, and the epistemic goods relevant for assessing their epistemic goodness. More specifically, we saw that the distinctive epistemic advantages of bureaucracies in terms of reliability and speed arise only if general conflicts between assigned social roles and personal interests are absent, if cost-efficient communication technology is available, and if the social system is located in a relatively stable environment. Furthermore, we saw that even under such favorable conditions, it remains a distinctive epistemic downside of bureaucracies that they obstruct the formation of understanding. We saw that this non-ideal systems-oriented social epistemology was able to explain both the epistemic failure in *FIASCO* as well as the epistemic success in *SUCCESS STORY*.

Let me end with some reflections on the significance of this result. On the one hand, this result is significant in a straightforward way because it allows us to better understand an important part of our lives and to identify countermeasures that can mitigate the distinctive epistemic downsides of bureaucracies. On the other hand, it allows us to make a more informed decision whether the favorable conditions under which bureaucracies yield good epistemic outcomes are realistically attainable and all things considered worth pursuing. Furthermore, the broader importance of this topic for social epistemology becomes apparent once we consider that it remains a widely discussed thesis that there is a general tendency towards bureaucratization. Weber (2015 [1921], 97) attributes this tendency to the superior efficiency of bureaucracies under market conditions. While his market-based explanation remains controversial, other mechanisms have been suggested that likewise predict a general tendency towards bureaucratization (Parkinson 1958; DiMaggio & Powell 1983; Graeber 2005, 149–205). To assess the correctness of these arguments lies beyond the scope of this article. However, if they are correct, then the epistemic assessment of bureaucracies undertaken here matters for the epistemic assessment of virtually *any* domain of society.²⁰

²⁰ I am grateful to John Greco for his written feedback on a previous draft of this article. This article also benefited from discussions at the 4th Social Epistemology Network Event (SENE4) at Georgetown University in September 2025, the APA Pacific Division Meeting in San Francisco in April 2025, and the Pluralistic Philosophy Workshop at Georgetown University in February 2023. In particular, I thank Andrius Bielskis, Gabriele Contessa, Dominick Cooper, Mona Fazeli, Will Fleisher, Carolina Flores, B. Huebner, Quill Kukla, Neil Levy, Elisa Reverman, Megan Ritz, Andy Sullivan, Yo-ran Yang, and two anonymous referees.

§ 7. References

- Arndt, B. G., Beasley, J. W., Watkinson, M. D., Temte, J. L., Tuan, W. J., Sinsky, C. A., & Gilchrist, V. J. (2017). Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations. *Annals of family medicine*, 15(5), 419–426.
- Beetham, D. (1987). *Bureaucracy*. Milton Keynes: Open University Press.
- Blackburn, S. (2018). *On truth*. Oxford University Press.
- Bratman, M. E. (1992). Shared cooperative activity. *Philosophical Review*, 101(2), 327–341.
- Burns, T., & Stalker, G. M. (1961). *The management of innovation*. London: Tavistock.
- Chin, G. J., & Wells, S. C. (1998). The ‘blue wall of silence’ as evidence of bias and motive to lie: A new approach to police perjury. *University of Pittsburgh Law Review*, 59, 233–299.
- Coase, R. (1937). The nature of the firm. *Economica*, 4(16), 386–405.
- Daston, L. (1994). Enlightenment calculations. *Critical Inquiry*, 21(1), 182–202.
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality. *American Sociological Review*, 48(2), 147–160.
- Dennett, D. C. (2017). *From bacteria to Bach and back: The evolution of minds*. New York: W. W. Norton & Company.
- Downs, A. (1966). *Bureaucratic structure and decisionmaking*. Santa Monica, CA: RAND Corporation.
- Elgin, C. Z. (2017). *True enough*. Cambridge, MA: The MIT Press.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Gilbert, M. (1989). *On social facts*. London / New York: Routledge.
- Goldman, A. I. (1988). *Epistemology and cognition*. Harvard University Press.
- Goldman, A. I. (2011). A guide to social epistemology. In A. I. Goldman & D. Whitcomb (Eds.), *Social epistemology: Essential readings* (pp. 11–37). Oxford University Press.
- Gouldner, A. W. (1954). *Patterns of industrial bureaucracy*. New York: The Free Press.
- Grattan-Guinness, I. (1990). Work for the hairdressers: The production of de Prony’s logarithmic and trigonometric tables. *Annals of the History of Computing*, 12(3), 177–185.
- Graeber, D. (2005). *The utopia of rules: On technology, stupidity, and the secret joys of bureaucracy*. Brooklyn: Melville House.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Greco, J. (2020). *The transmission of knowledge*. Cambridge University Press.
- Grier, D. A. (2005). *When computers were human*. Princeton University Press.
- Grimm, S. (2012). The value of understanding. *Philosophy Compass*, 7(2), 103–117.
- Herd, P., & Moynihan, D. P. (2018). *Administrative burden: Policymaking by other means*. New York: Russell Sage Foundation.
- Hodson, R., Martin, A. W., Lopez, S. H., & Roscigno, V. J. (2013). Rules don’t apply: Kafka’s insights on bureaucracy. *Organization*, 20(2), 256–278.
- Höpfl, H. (2006). Post-bureaucracy and Weber’s ‘modern’ bureaucrat. *Journal of Organizational Change Management*, 19, 8–21.
- Joukovskaia, A. (2023). Bureaucracy: The making of a buzzword. *Journal of the History of Ideas*, 84(4), 685–710.
- Kafka, F. (2009 [1926]). *The castle* (A. Bell, Trans.; R. Robertson, Ed.). Oxford University Press.
- Kolodny, A. (2020). How FDA failures contributed to the opioid crisis. *AMA Journal of Ethics*, 22(8), E743–E750.

- Kremers, P. (forthcoming). Non-ideal theory as ideology. *Inquiry*, 1–23.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Luhmann, N. (1972 [1964]). *Funktionen und Folgen formaler Organisation*. Berlin: Duncker & Humblot.
- Mayntz, R. (1965). Max Webers Idealtypus der Bürokratie und der Organisationssoziologie. In J. Fijalkowski (Ed.), *Politologie und Soziologie: Otto Stammer zum 65. Geburtstag* (pp. 91–100). Köln / Opladen: Westdeutscher Verlag.
- McKenna, R. (2023). *Non-ideal epistemology*. Oxford University Press.
- Merton, R. K. (1968 [1940]). Bureaucratic structure and personality. In *Social theory and social structure* (pp. 249–260). New York: The Free Press.
- Mills, C. W. (2005). ‘Ideal theory’ as ideology. *Hypatia*, 20(3), 165–184.
- National Human Genome Research Institute (2024). Human Genome Project fact sheet. *National Human Genome Research Institute*. <https://www.genome.gov/about-genomics/educational-resources/fact-sheets/human-genome-project>.
- Nguyen, C. T. (2021). The seductions of clarity. *Royal Institute of Philosophy Supplement*, 89, 227–255.
- Nijor, S., Rallis, G., Lad, N., & Gokcen, E. (2022). Patient Safety Issues From Information Overload in Electronic Medical Records. *Journal of patient safety*, 18(6), e999–e1003.
- Office of Strategic Services (1944). Simple sabotage field manual. *Homeland Security Digital Library*. <https://www.hsdl.org/c/abstract/?docid=750070>.
- Olson, M. (1965). *The logic of collective action: Public goods and the theory of groups*. Harvard University Press.
- Parkinson, C. N. (1958). *Parkinson’s Law, or The Pursuit of Progress*. London: Sidgwick & Jackson.
- Reverman, E. C. (2025). Testimonial compression. *Episteme*, 22, 797–817.
- Shannon, C. E., & Weaver, W. (1964 [1949]). *The mathematical theory of communication*. Urbana, IL: University of Illinois Press.
- Shapiro, S. (2014). Massively shared agency. In M. Vargas & G. Yaffe (Eds.), *Rational and social agency: The philosophy of Michael Bratman* (pp. 257–293). Oxford University Press.
- Shieber, J. (2013). Toward a truly social epistemology: Babbage, the division of mental labor, and the possibility of socially distributed warrant. *Philosophy and Phenomenological Research*, 86(2), 266–294.
- Simon, H. A. (1998 [1947]). *Administrative behavior*. New York: Macmillan.
- Simon, H. A., & March, J. G. (1993 [1958]). *Organizations*. Cambridge, MA: Wiley-Blackwell.
- Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, 6(1), 1–10.
- Smith, A. (1976 [1776]). *An inquiry into the nature and causes of the wealth of nations*. Oxford University Press.
- Stahl, T. (2024). What (if anything) is ideological about ideal theory? *European Journal of Political Theory*, 23(2), 135–158.
- Tullock, G. (1965). *The politics of bureaucracy*. Washington, DC: Public Affairs Press.
- Tuomela, R. (2007). *The philosophy of sociality: The shared point of view*. Oxford University Press.
- Weber, M. (1964 [1947]). *The theory of social and economic organization* (A. M. Henderson & T. Parsons, Trans.). New York: The Free Press.
- Weber, M. (1993 [1930]). *The Protestant ethic and the spirit of capitalism* (T. Parsons, Trans.). London / New York: Routledge.

- Weber, M. (1994 [1919]). The profession and vocation of politics. In P. Lassmann & R. Speirs (Eds.), *Weber: Political writings* (pp. 307–369). Cambridge University Press.
- Weber, M. (2015 [1921]). Bureaucracy. In T. Waters & D. Waters (Eds. & Trans.), *Weber's rationalism and modern society: New translations on politics, bureaucracy, and social stratification* (pp. 73–128). New York: Palgrave Macmillan.
- Williamson, O. E. (1967). Hierarchical control and optimum firm size. *Journal of Political Economy*, 75(2), 123–138.
- Zagzebski, L. (2003). The search for the source of epistemic good. *Metaphilosophy*, 34(1/2), 12–28.