

What Infrastructure Is For

Consider the domestic abuse survivor who cannot leave because every transaction is visible to the man who monitors her account. Her problem is not that the financial system is malfunctioning. It is functioning exactly as designed — recording everything, losing nothing, available to whoever has access. We have a standard method for evaluating technologies, and it is the wrong method for the most important class of them. We ask: what are the average effects on average users in ordinary conditions? But there is a class of infrastructure whose value is not located at the mean — whose value is realized precisely at the extremes, by people for whom the architecture of permission has become immediately consequential rather than abstractly political. Evaluate that infrastructure by its average effects and you will reach the wrong conclusion every time. The cashless question is not primarily a question about payments. It is a question about whether we understand what infrastructure is for — and specifically, whether we understand why certain infrastructures need to remain available to the person in the worst situation, for whom the gap between legal constraint and architectural impossibility is the difference between survival and its absence.

There is a categorical difference between a law that prohibits anonymous transactions and a world in which anonymous transactions cannot technically occur. The first is a legal constraint: it can be challenged, defied at cost, repealed, or broken by someone willing to accept the consequences. It creates friction rather than impossibility. The second is architectural: it does not prohibit anything. It simply makes certain acts cease to exist as options. The freedom it eliminates is not restricted. It is abolished. And once abolished, cash infrastructure cannot be restored at will: payment networks consolidate around dominant platforms through strong network effects, financial institutions build compliance architecture on the assumption of full transaction visibility, and the political economy of reintroduction faces an industry with structural incentives to resist it. This is not permanence in any absolute sense — but it is path dependency deep enough that "we can always reverse it" is not a serious answer. The surveillance objection to cashlessness, as usually stated, misses this: it worries about bad actors misusing data, about chilling effects, about creeping normalization. These are mean-case

concerns. The deeper problem is that architectural elimination forecloses a category of freedom rather than constraining its exercise — and that is a different problem entirely.

The optimist's case for cashlessness is real and should not be dismissed. Digital payment infrastructure has done more for financial inclusion in a decade than physical banking managed in a century; tax evasion facilitated by cash runs into the trillions annually; criminal networks depend on its untraceability. The mean-case analysis is not wrong. It is, for the purposes it addresses, correct. The problem is not the analysis. It is the assumption that the mean case is the right place to look.

Philip Pettit's republicanism and Albert Hirschman's exit theory arrive at the same place from different directions, and that convergence is important. For Pettit, freedom is non-domination: the condition of not being subject to another's arbitrary will, regardless of whether that will is currently being exercised. A slave whose master is currently benevolent is not free — the domination is real even when its exercise is suspended. What cash provides is not primarily the freedom to transact anonymously. It is the structural guarantee that one *could* — the architecture of exit that prevents the relationship between individuals and financial institutions from becoming one of domination, even when those institutions are behaving well. Hirschman's framework makes the political mechanism precise: institutions are disciplined by the credible possibility that participants can exit. Voice — complaint, advocacy, political pressure — benefits those already powerful within the institution. Exit is available to anyone who can leave. Cash is one of the last mechanisms by which an ordinary person can exit the formal financial system entirely. The domestic abuse survivor who needs to move money without her abuser knowing is not making a theoretical claim. She is navigating a practical architecture, and whether exit is available to her is the question her situation actually poses. Pettit and Hirschman are not making separate points. They are describing the same structural fact from two angles: that the guarantee of exit is what converts institutional benevolence from a gift into an obligation. Remove the guarantee, and benevolence has no anchor. Trust that has no alternative is not trust. It is dependency wearing trust's vocabulary.

The strongest objection insists that this romanticizes a tool whose primary beneficiaries are not abuse survivors and dissidents but tax evaders, drug traffickers, and the already-comfortable who

prefer opacity for reasons of convenience rather than survival. The domestic violence literature cuts further still: abusers use cash to maintain financial control precisely because it is untraceable, keeping partners from monitoring what is concealed. These harms are concentrated and severe, and the argument from tails must answer them honestly.

The answer is this: the harms are real, and they do not decide the question — because the question is not amenable to the method by which they are weighed. In social choice theory, a lexicographic preference ordering is one in which a higher-order consideration must be fully satisfied before lower-order considerations are evaluated at all. Structural freedom functions as a lexical constraint of exactly this kind — not because freedom is intrinsically weightier than welfare, but because welfare calculations presuppose agents capable of contesting their outcomes, and that capacity requires exit. An agent who cannot act outside institutional permission cannot meaningfully challenge the terms on which welfare is distributed to her. Structural freedom is therefore prior to welfare analysis in the specific sense that it is the condition under which welfare analysis has democratic legitimacy rather than merely technical validity. A policy that improves aggregate welfare while eliminating that structural condition has not made a trade visible in the welfare calculation. It has removed the foundation on which the legitimacy of that calculation rests. When we design away anonymous transaction to constrain its misuse, we achieve real gains at the mean and make a separate, incommensurable loss at the tail — a loss invisible to the method that makes the gains look decisive.

The mean-case method is not a bias or an oversight. It is the natural product of designing systems for the people who use them most, in the conditions that obtain most often. What it cannot see is the person for whom ordinary conditions have already failed. Whether we protect the structural guarantee of exit for that person, even when the aggregate calculus runs against it, is not a question about payment technology. It is a question about what a society believes freedom is for.