Systemic Arrestability: Operator Protection in Machine Safety Law When Machine Speed Exceeds Human Reaction

Giovanni Nardacci
Human Flag Association, Bellinzona, Switzerland

Abstract

Machine safety law rests on a foundational principle: the protection of the operator. Across legal systems, safety regulations require that machines be designed and operated so that dangerous situations can be stopped before injury occurs. This protection is not a new legal objective but an obligation already embedded in existing law — the stop-function and emergency-stop requirements of Directive 2006/42/EC (and the forthcoming Regulation (EU) 2023/1230, which expressly addresses machines with autonomous behaviour), the product-safety duties of Switzerland's LSPro (RS 930.11), and the operator-protection requirements of OSHA (29 CFR § 1910.212).

This paper argues that such protection rests on an implicit temporal assumption: that the operator retains a meaningful opportunity to perceive a developing hazard, evaluate it, and intervene before harm occurs. Traditional machinery generally satisfies this assumption; it can fail whenever a machine produces operationally significant effects faster than human reaction allows. In such circumstances the operator's legal protection formally remains in place, while the practical conditions for exercising it disappear. The resulting problem is not regulatory absence but legal effectiveness: existing obligations remain valid yet can no longer be realised through human intervention alone.

To address this gap, the paper introduces the concept of *systemic arrestability*: a system is systemically arrestable when its capacity to halt a hazardous action does not depend on a human operator perceiving, evaluating, and intervening within the time window of that action. The concept creates no new legal obligation; it preserves the practical effectiveness of obligations that already exist. Through a comparative analysis of the European, Swiss, and United States machine-safety frameworks, the paper proposes systemic arrestability as an interpretative standard for maintaining the protective purpose of existing safety law under conditions in which machine speed exceeds human reaction capacity.

Keywords: operator safety, machine safety, stop function, emergency stop, human reaction time, legal effectiveness, systemic arrestability, Directive 2006/42/EC, Regulation (EU) 2023/1230, LSPro RS 930.11, OSHA 29 CFR § 1910.212.

1. Introduction

Machine safety law begins with the operator. Long before contemporary debates on artificial intelligence and autonomy, legal systems on both sides of the Atlantic converged on a simple structural principle: a machine must be designed and operated so that the person using it can stop a dangerous situation before it causes harm. This principle is not aspirational. It is codified in binding norms: the stop-function and emergency-stop requirements of Directive 2006/42/EC (EHSR 1.2.4.1 and 1.2.4.3), the general duty of safe design under EHSR 1.1.2, the product-safety obligations of Switzerland's Federal Act on Product Safety (LSPro, RS 930.11, arts. 3–4), and the machine-guarding requirements of the United States Occupational Safety and Health

Administration (29 CFR § 1910.212(a)(1)), which expressly protect "the operator and other employees in the machine area."[1]

These provisions differ in structure, scope, and enforcement. Yet they share an assumption so fundamental that it has rarely needed to be stated: that between the onset of a hazardous machine action and the occurrence of harm, there exists a time window within which a human being can perceive the hazard, evaluate it, and intervene. The emergency-stop button presupposes a hand that reaches it in time. The stop function presupposes an operator who recognises the need to use it. The duty to guard presupposes a danger that develops at a speed commensurate with human reaction. This assumption — referred to here as the *temporal assumption* of machine safety law — was reasonable for the machinery these norms were written to govern, and it remains satisfied across most of the industrial landscape.

It is not, however, a law of nature. A growing class of machines can produce *operationally significant effects* — defined here as effects whose occurrence materially impairs the possibility of preventing harm through subsequent intervention; irreversibility and injuriousness are the typical indicators of such effects, not the definition itself — within time windows shorter than documented human reaction capacity. Where this occurs, a structural inversion takes place: the operator's legal protection formally remains in force, while the practical conditions for exercising it cease to exist. The operator is not negligent, untrained, or inattentive; the operator is objectively exposed to a risk that no degree of skill or vigilance can control, because the decisive interval has closed before human cognition can open it. Machines exhibiting autonomous or self-evolving behaviour are a prominent example of this class — and the European legislator has recognised as much in Regulation (EU) 2023/1230, applicable from January 2027, which requires for the control systems of such machinery that it be possible at all times to correct the machinery in order to maintain its inherent safety. But autonomy is an instance of the problem, not its foundation. The problem is temporal, and it arises wherever machine speed exceeds human reaction, whatever the underlying technology.

This paper makes three claims. First, that the temporal assumption is genuinely embedded in existing machine safety law, identifiable in the text and structure of the European, Swiss, and American frameworks (Section 2). Second, that when the assumption fails, the result is not a regulatory gap but an *effectiveness* gap: obligations remain valid while becoming impossible to discharge through human intervention alone — a distinction with significant consequences, because it means the solution lies in interpretation and application of existing law, not in new legislation (Section 3). Third, that this effectiveness gap can be closed by an interpretative criterion the paper terms *systemic arrestability*: a system is systemically arrestable when its capacity to halt a hazardous action does not depend on a human operator perceiving, evaluating, and intervening within the time window of that action (Section 4). The criterion is legal-functional, not technical: it prescribes no particular engineering solution, and — as Section 4 also shows — it creates no obligation that does not already exist. It restates, for a specific class of machines, what the duty of safe design has always required.

[1]Directive 2006/42/EC of the European Parliament and of the Council of 17 May 2006 on machinery, OJ L 157, 9.6.2006, Annex I (Essential Health and Safety Requirements). Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery, OJ L 165, 29.6.2023, applicable from 20 January 2027 (Art. 54), Annex III.

Methodologically, the paper is doctrinal and comparative. It does not propose a technical safety standard; it reconstructs an implicit legal premise shared by three machine-safety regimes and derives from it an interpretative criterion. Engineering enters the analysis only as evidence of feasibility, never as the source of the legal argument.

The analysis is confined to civil and dual-use machinery, the domain in which the legal anchors examined here apply. Military systems fall outside civil product-safety law and are noted only as an open question (Section 5). The protection of third parties is not treated separately: the same mechanism that protects the operator necessarily protects those in the machine's vicinity, a point the American regulation makes textually explicit.[2]

2. The Temporal Assumption in Machine Safety Law

This section examines the three legal frameworks that anchor the analysis — European Union machinery law, Swiss product-safety law, and United States occupational safety regulation — and shows that each, in its own structure, presupposes a human operator who has time to act. The point of the comparison is not similarity of drafting but convergence of architecture: three legal orders, with different traditions and enforcement mechanisms, protect the operator through devices and duties that are intelligible only if a window for human intervention exists.

2.1 European Union: Directive 2006/42/EC and Regulation (EU) 2023/1230

The Machinery Directive operates through Essential Health and Safety Requirements (EHSRs) that manufacturers must satisfy before placing machinery on the market. Three are central here.

EHSR 1.1.2 establishes the principle of safety integration: machinery must be designed and constructed so that it can be operated, adjusted and maintained without putting persons at risk, with risks eliminated or reduced by design as the first and preferred measure, protective measures second, and information to users only as a residual instrument. The hierarchy itself is significant: the law already prefers solutions that do not depend on human conduct. Yet the Directive's specific stopping provisions reveal where its drafters located the ordinary case.

EHSR 1.2.4.1 requires that machinery be fitted with a control device whereby it can be brought safely to a complete stop. EHSR 1.2.4.3 requires one or more emergency stop devices "to enable actual or impending danger to be averted." Both provisions are instruments placed *in human hands*. A stop control presupposes someone who decides to stop the machine; an emergency stop presupposes someone who recognises an "actual or impending danger" and physically reaches the device while the danger is still impending rather than consummated. The temporal assumption is not stated because it did not need to be: for the machinery of 2006, the interval between "impending" and "actual" was a human-scale interval.

Nor is the assumption merely structural; it surfaces in the text. EHSR 1.2.2 requires that control devices be positioned in such a way as to be safely operated "without hesitation or loss of time and without ambiguity"; EHSR 1.2.4.3 requires the emergency stop to halt the hazardous process "as quickly as possible." The Commission's own interpretive material on the emergency-stop

[2]All normative citations in this paper have been verified against the official texts (EUR-Lex consolidated text CELEX 02023R1230; ecfr.gov; fedlex.admin.ch / admin.ch) as of June 2026. They remain subject to confirmation by qualified counsel prior to formal submission.

requirement is explicit about the reason: in an emergency situation, a split-second reaction may be crucial. The law speaks of time because it assumes a race the human can still win.[3]

The Union judicature has confirmed both the hierarchy and the method. In *CSF v Commission*, the General Court held that a specific safety requirement of Annex I must be interpreted in the light of the General Principles and of the principles of safety integration set out in EHSR 1.1.2; that the Directive "does not merely impose an obligation on manufacturers to warn their customers against the risks" but also obliges them to eliminate or reduce those risks as far as possible at the time of design and construction; and that once a risk is ascertained, "the fact that the user has received prior notice of the existence of the risk is in itself irrelevant, given both the hierarchy of the prevention and information requirements" imposed on manufacturers. Two holdings matter for the present argument: information addressed to the human cannot substitute for design, and specific provisions are to be read in the light of the general duty — the precise interpretative operation that Section 4 will propose for the temporal dimension.[4]

The European legislator has since acknowledged that this can no longer be taken for granted. Regulation (EU) 2023/1230, which replaces the Directive with effect from 20 January 2027, provides that control systems of machinery with fully or partially self-evolving behaviour or logic, designed to operate with varying levels of autonomy, shall be designed and constructed in such a way that "it shall be possible at all times to correct the machinery or related product in order to maintain its inherent safety." The same provision requires that such systems not cause the machinery to perform actions beyond its defined task and movement space, and the Regulation's general principles extend the manufacturer's risk assessment to hazards arising as an intended evolution of the machinery's self-evolving behaviour or logic.[5]

Three features of the correction requirement matter for the present argument. First, it confirms that the legislator regards continuous correctability as already inherent in the safety of such machines — the Regulation speaks the language of *maintaining* inherent safety, not of adding a new layer to it. Second, it is located within the requirements on control systems (Annex III, section 1.2.1), that is, within design, not within operator instructions. Third, it is category-specific: it attaches to self-evolving behaviour as such, whereas the temporal problem examined in this paper attaches to speed, and can arise in machinery that exhibits no autonomy at all. The Regulation is therefore best read as the legislator's first explicit response to a problem that is wider than its trigger.

2.2 Switzerland: the Federal Act on Product Safety (LSPro, RS 930.11)

Swiss law approaches the same architecture from the product side. Article 3(1) LSPro permits a product to be placed on the market only if, used normally or in reasonably foreseeable conditions, it exposes the safety and health of users *and third parties* to no danger, or only to minimal danger. Article 4 empowers the Federal Council to establish the essential health and safety requirements for specific product categories in sectoral legislation — the channel through which machine-specific requirements aligned with the European EHSRs enter Swiss law. The duty is the manufacturer's

[3]European Commission, Machinery Working Group document MD WG 2010-03 EN rev. 2 (approved 5–6 November 2014), on emergency stop devices under Annex I, section 1.2.4.3 of Directive 2006/42/EC, quoting the Guide to application of the Machinery Directive: "in an emergency situation, a split-second reaction may be crucial."

[4]Case T-337/13, *CSF Srl v European Commission*, judgment of the General Court (Third Chamber) of 15 July 2015, in particular the holdings summarised at paras 64–70 and 83–84.

[5]Regulation (EU) 2023/1230, Annex III, section 1.2.1, third subparagraph, point (c).

and it is anterior to use: safety must exist in the product as placed on the market, not be supplied afterwards by the user's prudence.[6]

Two features of the Swiss framework sharpen the argument. First, the express inclusion of third parties in Article 3(1) confirms that operator protection and bystander protection are legally one mechanism, not two: the product that is safe for its user is, by the same design, safe for those around it. Second, Articles 16 and 17 LSPro attach criminal liability to violations of the product-safety duties. This raises the stakes of the temporal question. What the criminal provisions sanction is the placing on the market of a product that does not meet the safety requirements; if a product's safety case rests on operator intervention that the time window excludes, the product arguably never satisfied Article 3(1) at the moment it was placed on the market. The criminal provisions thus convert the temporal assumption from an interpretative observation into a matter of practical exposure: the Swiss manufacturer must ask not whether a stop device exists, but whether it can ever be used in time.

2.3 United States: OSHA 29 CFR § 1910.212

The American anchor is the most textually direct. Section 1910.212(a)(1) requires that "one or more methods of machine guarding shall be provided to protect the operator and other employees in the machine area from hazards such as those created by point of operation, ingoing nip points, rotating parts, flying chips and sparks." Three aspects deserve emphasis.[7]

First, the provision names its beneficiary: the operator, first in the textual order, with other employees in the machine area alongside. The dual protection that Swiss law states as "users and third parties" appears here as a single occupational community gathered around the machine.

Second, the regulation is method-neutral but outcome-fixed: it does not prescribe a particular guard — its own examples range from barrier guards to two-hand tripping devices to electronic safety devices — but it requires that *some* method actually protect. A method that depends on the operator reacting faster than the hazard develops does not protect the operator; it assigns the operator the task of protecting himself, which is precisely what the provision exists to displace. The point is made textually explicit one paragraph later: under § 1910.212(a)(3), point-of-operation guarding shall be "so designed and constructed as to prevent the operator from having any part of his body in the danger zone during the operating cycle." Where reaction cannot be trusted, the regulation does not demand better reaction; it designs the operator out of the danger zone altogether. The Review Commission's case law draws the same line: it has long held that work rules governing the operator's conduct are not a method of guarding contemplated by § 1910.212(a)(1); adequacy of guarding is assessed against the manner in which the machine functions and how it is operated by the employees; and exposure is established wherever it is reasonably predictable that employees have been, are, or will be in the zone of danger. Reliance on operator conduct, in other words, is not guarding — in American law as a matter of decided cases.[8]

---

[6]Loi fédérale sur la sécurité des produits du 12 juin 2009 (LSPro), RS 930.11 / Legge federale sulla sicurezza dei prodotti (LSPro), RS 930.11, in force since 1 July 2010.

[7]29 CFR § 1910.212 (General requirements for all machines), Subpart O – Machinery and Machine Guarding.

[8]*Akron Brick & Block Co.*, 3 BNA OSHC 1876, 1877–78 (No. 4859, 1976); *Rockwell Int'l Corp.*, 9 BNA OSHC 1092, 1097–98 (No. 12470, 1980); *Fabricated Metal Products, Inc.*, 18 BNA OSHC 1072, 1073–74 (No. 93-1853, 1997). For a recent application of these precedents, see *Aerospace Testing Alliance*, OSHRC Docket No. 16-1167 (Sept. 21, 2020).

Third, the hazards enumerated — nip points, rotating parts, flying chips — are instructive. They are fast hazards. American machine-guarding law was built around dangers that already exceeded human reaction time at the point of operation, and its answer was never "react faster": it was physical guarding that makes the dangerous interval humanly irrelevant. In this sense, § 1910.212 contains the seed of the argument this paper generalises: where reaction is impossible, the law has always demanded design.

2.4 Convergence

Read together, the three frameworks display a common structure. Each imposes the safety duty before and above the operator's own conduct (safety integration; market-placement conditions; mandatory guarding). Each provides stopping or guarding instruments whose ordinary operation presupposes human perception and reaction — and, in the European instruments, says so in terms, requiring operation "without hesitation or loss of time" and stopping "as quickly as possible." And each contains, at least in embryo, the recognition that where human reaction cannot bridge the dangerous interval, the duty migrates to design — explicitly in EHSR 1.1.2's hierarchy and § 1910.212's guarding logic, prospectively in Regulation 2023/1230's correction requirement, and structurally in the LSPro's anterior product duty backed by criminal liability.

The temporal assumption, in short, is real, identifiable, and load-bearing in all three systems — and in the European texts it is not even implicit.

2.5 Comparative confirmation beyond the Euro-Atlantic frameworks

The three frameworks examined above suffice to establish the argument. A reader may nonetheless ask whether the structure they share is a feature of machine safety law as such, or a construction peculiar to Western regulatory traditions. This subsection addresses that question at the level of confirmation, not of full comparative treatment, and with a methodological caution stated openly: the primary sources of the systems considered here are in Russian, Chinese, and Hebrew, and the analysis relies on published translations of binding texts and on the documented adoption of international standards. The claim made is correspondingly architectural — that the same protective structure recurs — not exegetical. Accordingly, this subsection is offered as corroborative, not constitutive, evidence of the paper's thesis: the argument stands on the three frameworks examined above, and would stand without what follows.

The Eurasian Economic Union regulates machinery through Technical Regulation TR CU 010/2011 "On safety of machinery and equipment", a binding instrument in force since 2013 across Russia, Belarus, Kazakhstan and the other member states. Its architecture is recognisable at once: minimum safety requirements imposed on the manufacturer across the entire lifecycle, from design onward; coordination with — though not identity to — Directive 2006/42/EC; and stopping duties stated in terms. Control systems of machinery must be equipped with emergency braking and emergency stop means where their use can reduce or prevent danger; the emergency stop device must remain in the stop position until deliberately reset, and its reset must not by itself restart the machine. The requirement that safety be designed into the product before the operator ever acts — and that a stopping capacity be present wherever it can avert danger — is, in the EAEU, statutory law.[9]

[9]Technical Regulation of the Customs Union TR CU 010/2011 "On safety of machinery and equipment", adopted by Decision No. 823 of the Customs Union Commission of 18 October 2011, in force since 15 February

The People's Republic of China reaches the same architecture through a different vector: the adoption of international standards as national ones. GB/T 15706 is the Chinese adoption of ISO 12100, the foundational standard on safe design and risk reduction; GB 5226.1 is the adoption of IEC 60204-1, which carries the emergency-stop requirements; GB/T 16855.1 adopts ISO 13849-1 on safety-related parts of control systems. Each is designated "IDT" — identical — in the national catalogue: the normative material is the same, recepted. The design hierarchy, the stop function, and the guarding logic examined in Sections 2.1–2.3 are therefore not translated approximations in Chinese machine-safety practice; they are textually the same provisions.[10]

Israel's framework, anchored in the Work Safety Ordinance (New Version) of 1970, carries forward the British factories-legislation tradition from which American machine-guarding law also descends — the duty to securely fence dangerous parts of machinery — while the mandatory standards declared under Israeli law are adoptions of the corresponding IEC and ISO standards. The Israeli case is cited here with deliberate restraint, at the level of architecture, pending primary-source verification; its interest lies in the fact that an ecosystem dense with autonomous and safety-critical systems rests its machine-safety law on the same two pillars — fencing duties of British descent and adopted international standards — found elsewhere in this survey.[11]

The significance of this confirmation is precise. The temporal assumption appears not merely across three Western systems, but across legally and politically diverse machine-safety regimes — statutory in the European Union, Switzerland, the United States and the Eurasian Economic Union, and carried into Chinese and Israeli law through the identical adoption of the international standards that embody it. The role of those standards must be stated carefully, in line with Section 4.3: as voluntary instruments they bind no court, and nothing in this paper derives an obligation from them. Their evidentiary role here is different — they demonstrate that the protective architecture resting on the temporal assumption is the common engineering and regulatory inheritance of systems that agree on little else. What none of these systems yet states in general terms is the criterion that governs when the assumption fails and what its failure requires. That is the task of Sections 3 and 4.

2013, binding throughout the Eurasian Economic Union. Quotations follow published English translations of the official Russian text; the original is authoritative.

[10]GB/T 15706 (Safety of machinery — General principles for design — Risk assessment and risk reduction), the Chinese national adoption of ISO 12100, designated "IDT" (identical); GB 5226.1 (Electrical safety of machinery — Electrical equipment of machines — Part 1), the identical adoption of IEC 60204-1, which contains the emergency-stop requirements; GB/T 16855.1, the identical adoption of ISO 13849-1 (safety-related parts of control systems).

[11]Work Safety Ordinance (New Version), 5730-1970 (Israel), carrying forward the machinery-fencing duties of the British factories-legislation tradition; mandatory Israeli standards declared under the Standards Law are adoptions of the corresponding IEC/ISO standards by the Standards Institution of Israel. The Israeli framework is cited at the level of architecture only; primary-source verification in Hebrew remains outstanding. Within the legal heritage to which Israel's Foundations of Law Act, 5740-1980, assigns a residual interpretive role — directing courts, in cases of lacuna, to the principles of "freedom, justice, equity and peace of Israel's heritage", since 2018 expressly including Jewish law (*ha-mishpat ha-ivri*) — the duty of safe design is ancient: Deuteronomy 22:8 commands the builder of a new house to construct a parapet (*ma'akeh*) on its roof, an obligation that attaches to construction, before any negligence in use can occur. The reference is not offered as a source of contemporary machine-safety obligations, but as a historical illustration — corroborative, like this subsection as a whole — of the design-based allocation of safety duties within a legal tradition that Israeli law continues to recognise as interpretatively relevant. The structure of the safety duty examined throughout this paper — lodged in design rather than in conduct — thus long predates the frameworks that now codify it.

3. When the Assumption Fails

3.1 The human side of the window: reaction as a chain, not an instant

The temporal assumption fails as a matter of fact, not of doctrine, and its failure can be described with precision. Human intervention on a machine is not a single act but a chain: the hazard must produce a perceptible signal; the signal must be perceived; the perception must be evaluated and a decision formed; the decision must be executed as a physical act on a control; and the control must take effect on the machine. Each link consumes time, and the times compound.

The empirical literature on human performance gives orders of magnitude rather than fixed constants, and this paper needs nothing more precise. Simple reaction to an expected stimulus occupies roughly a fifth of a second; recognition and choice among alternatives lengthen the interval substantially; and realistic operational conditions — an unexpected hazard, an ambiguous signal, a control at arm's length rather than under the finger — bring total response times into the range of one to several seconds. These figures describe trained, attentive, unimpaired adults. They are not deficiencies; they are the physiology the law's protection was built for, and no instruction, training, or warning can shorten them below their floor.[12]

Two consequences follow. First, the window is irreducible: it can be narrowed by design (better signals, closer controls) but never closed, because its core components are neurological. Second, the window is *knowable in advance*: a manufacturer can estimate, for a given machine and operating position, the minimum credible human response time. The temporal assumption is therefore not an imponderable; it is a parameter that design can be measured against.

3.2 The machine side of the window: operationally significant effects

On the other side of the comparison stands the speed at which a machine produces *operationally significant effects* — effects whose occurrence materially impairs the possibility of preventing harm through subsequent intervention. The definition is functional rather than physical: it is anchored to the protective purpose of the safety provisions, not to the characteristics of the event. Irreversibility and injuriousness are the typical indicators that such an effect has occurred, but they are evidence of the definition, not the definition itself. It is likewise effect-based rather than technology-based: it does not ask how the machine generates its action (mechanically, electronically, algorithmically) but what the action accomplishes within a given interval.

This is where the analysis must be kept on its proper footing. The question is never whether a machine is "fast" in the abstract — many machines perform fast motions in guarded spaces, and the guarding is precisely what makes their speed legally irrelevant. The question is whether the machine can produce operationally significant effects, in a space where persons are exposed, faster than the minimum credible human response time for that configuration. When it can, the comparison of the two windows yields a determinate result: the chain of human intervention cannot complete before the effect occurs. The emergency stop exists; the hand cannot reach it. The stop function exists; the decision to use it cannot form in time.

[12]The most extensively studied operational reaction chain is driver perception-brake response. Green's methodological review reports approximately 0.70–0.75 s for fully expected signals, about 1.25 s for unexpected but common events, and roughly 1.5 s for surprise events: M. Green, "'How Long Does It Take to Stop?' Methodological Analysis of Driver Perception-Brake Times", *Transportation Human Factors* 2(3) (2000) 195–216. See also H. Summala, "Brake Reaction Times and Driver Behavior Analysis", ibid., 217–226, reporting on-road braking latencies of 1.0–1.3 s in fairly urgent situations.

3.3 The inversion: protection in force, protection unusable

What occurs at that point is best described as a structural inversion. Every element of the operator's legal protection remains formally intact — the stop devices are installed, conformity is declared, the documentation is complete. Yet the protective architecture has silently changed its nature: instruments designed to *give the operator control over danger* have become instruments that *presuppose a control the operator cannot exercise*. The machine is formally compliant and functionally unguarded.

It must be stated plainly that this is a condition of fact and not a judgment about persons. The operator in this position is not negligent, untrained, slow, or inattentive; the most capable operator conceivable stands in exactly the same position, because the interval has closed before cognition can open it. The operator is objectively exposed to a risk that no quality of performance can control. Equally, the manufacturer in this position has not necessarily violated any specific provision examined so far; it has done something subtler — it has discharged a substantive duty through instruments that cannot perform their substantive function. The wrong, if it ripens into one, lies in reliance on human intervention as the safety mechanism in conditions where human intervention is physically excluded.

3.4 Autonomy as instance, speed as essence

Machines with autonomous or self-evolving behaviour are the most discussed members of this class, and Regulation (EU) 2023/1230 confirms that the European legislator sees them as raising exactly this problem. But the analysis above makes clear why autonomy is an instance and not the essence. Autonomy aggravates the temporal problem in characteristic ways — it can originate hazardous action without a prior human command (so the perceptible signal arrives later in the chain), and it can vary its behaviour (so evaluation takes longer because the operator cannot rely on a learned pattern). Both aggravations, however, act on the same variable: they widen the gap between machine effect time and human response time. A non-autonomous system that produces irreversible effects in milliseconds presents the identical legal problem; an autonomous system slow enough to be interrupted does not present it at all.

Anchoring the analysis to speed rather than autonomy has a further advantage: it keeps the criterion stable across technological change. Definitions of autonomy age quickly; the comparison between two time windows does not.

3.5 The gap, restated

The result of Sections 2 and 3 can now be stated in one sentence: existing machine safety law protects the operator through duties and devices that presuppose a humanly bridgeable interval between hazard and harm, and a definable class of machines — identifiable in advance, by comparison of measurable time windows — eliminates that interval while leaving the duties and devices formally in place. This is an effectiveness gap, not a regulatory one: nothing needs to be added to the law, but something needs to be said about how the law applies where its central assumption fails. Section 4 proposes the criterion.

4. Systemic Arrestability

4.1 Definition and operational criteria

This paper proposes the following criterion. A system is *systemically arrestable* when its capacity to halt a hazardous action does not depend on a human operator perceiving, evaluating, and intervening within the time window of that action.

The definition resolves into two operational criteria, which are cumulative. First, the system must *allow time*: it must be designed so that, wherever feasible, a humanly bridgeable interval exists between the onset of a hazardous action and the occurrence of operationally significant effects. Second, where such an interval cannot exist, the system must *remain disableable within the actual window*: the capacity to halt the hazardous action must reside in a mechanism — physical, logical, or architectural — whose own response time fits inside the window that human response cannot. The first criterion preserves human intervention where it is possible; the second replaces dependence on human intervention where it is not. Together they restate, in temporal terms, what a stop function is for.

Three negative clarifications complete the definition. Systemic arrestability is not a prescribed technical solution: it does not mandate watchdogs, interlocks, hardware kill-switches, or any particular architecture, and it is satisfied by any mechanism that meets the temporal test. It is not a human-exclusion principle: it does not remove the operator's stop devices or diminish their role where the window permits their use; it ensures only that the *last* layer of stoppability does not depend on them. And it is not a standard of perfection: it requires that the halting capacity fit the window of the hazardous action, not that every conceivable failure be preventable.

4.2 The criterion is already practised

If systemic arrestability sounded utopian, it would deserve suspicion. It is, instead, ordinary industrial practice. Its purest instances are the interlock and the scram. A machine-tool interlock cuts power on the opening of a guard, without asking whether the operator noticed the danger; a reactor protection system scrams on parameter excursions precisely because no control-room crew can be trusted to react within the relevant interval. In both cases the halting capacity is lodged in the system itself, and human perception is removed from the critical path — which is exactly what the criterion requires. These belong to a broader family of safety functions engineered to act below the threshold of human reaction, from anti-lock braking to airbag deployment; the family demonstrates how routine human-independent response has become, but the interlock and the scram are the canonical cases, because what they do is precisely *arrest*.

These examples carry two arguments at once. They show feasibility: the criterion demands nothing the state of the art has not delivered routinely, across industries, for decades. And they show normality of cost: systemic arrestability has not priced machine tools or power plants out of existence; it has become invisible infrastructure. A criterion already satisfied by the humble interlock cannot be dismissed as a burden invented for emerging technology.

4.3 The objection: is this not already implicit in existing law?

A prepared reader will raise the strongest objection at this point, and it deserves to be stated at full strength. Existing law, the objection runs, already contains everything this paper proposes. EHSR

1.1.2 imposes safety integration: risks must be eliminated by design first, and only residually managed through the operator. The technical standards that give the EHSRs operational content — the functional-safety families governing control systems and protective devices — already define safety functions, performance levels, and response times that do not depend on human reaction. OSHA's guarding logic, as Section 2.3 showed, has always demanded design where reaction is impossible. If the duty already migrates to design whenever human intervention fails, what does *systemic arrestability* add beyond a name?

The objection is correct in everything it asserts and mistaken in what it concludes, and the difference is precisely the gap this paper has been describing. What existing law contains is the *material* from which the criterion is built; what it lacks is the criterion *stated as law*. The migration of the duty from operator to design currently lives in three places, none of which is a generally applicable legal proposition: in a design hierarchy (EHSR 1.1.2) whose temporal dimension is implicit and must be reconstructed interpretatively; in technical standards, which are voluntary instruments creating presumptions of conformity, addressed to engineers, revisable by committees, and binding on no court; and in category-specific provisions — guarding rules built around mechanical hazards, a correction requirement attached to self-evolving machinery — each of which captures one region of the problem without stating its principle.

The consequence of this dispersion is practical, not aesthetic. A manufacturer assessing a fast system today finds no single legal proposition that tells it the decisive question — *does your halting capacity fit inside the window of your hazardous action, independently of human reaction?* — and a court assessing an accident must assemble that question from fragments of a hierarchy, an annex, and a standard it is not bound to apply. Dispersed law is weak law in a precise sense: it permits formal compliance to substitute for functional protection, which is exactly the inversion Section 3.3 described. Systemic arrestability adds no obligation to this material; it consolidates the material into a criterion. It performs for the temporal dimension of machine safety the operation that legal concepts classically perform: it makes explicit, general, and applicable what was implicit, dispersed, and category-bound.

This is why the criterion is proposed as an *interpretative standard* and not as draft legislation. It is a way of reading duties that already bind — a canon for giving the stop-function, product-safety, and guarding provisions their effective meaning under temporal conditions their drafters did not need to contemplate. The operation has judicial precedent: it is the reading method the General Court applied in *CSF v Commission*, interpreting a specific requirement of Annex I in the light of the general principles of safety integration — and holding, in consequence, that information addressed to the operator cannot substitute for design. Nothing in the criterion requires legislative action, and nothing in it would be displaced by legislative action: Regulation 2023/1230's correction requirement is, on this reading, the first statutory appearance of a principle that the criterion states in general form.

4.4 What the criterion does in operation

Stated as an interpretative standard, systemic arrestability does determinate work at three junctures. For the manufacturer, it converts the temporal assumption from background into checklist: the comparison of windows described in Section 3 — minimum credible human response time against time-to-operationally-significant-effect — becomes part of risk assessment, and reliance on operator intervention becomes a design choice that must be justified by the window, not assumed.

For the assessor of conformity, it supplies the question that the fragmented material does not state in one place. And for the adjudicator, it gives content to duties whose breach is otherwise hard to articulate: a machine that depended on humanly impossible intervention did not satisfy the duty of safe design, however complete its documentation — not because a new duty has been created, but because the existing duty, read against the actual window, was never discharged.

The criterion's effect-based formulation also keeps it stable. Because it attaches to the comparison of time windows rather than to any technology, it does not age with definitions of autonomy, does not require amendment as architectures change, and applies identically to the millisecond mechanical hazard and the self-evolving controller. It is, in this respect, a deliberately old-fashioned legal object: a principle.

5. Implications and Open Questions

5.1 Third parties: one mechanism, not two

Nothing in this paper has argued for the protection of bystanders, and nothing needed to. The legal anchors themselves close the point: the LSPro extends the product-safety duty to "users and third parties" in the same breath; § 1910.212(a)(1) protects "the operator and other employees in the machine area" in a single clause. Protection of the person at the machine and protection of the persons around it are not two obligations requiring two mechanisms; they are one design outcome viewed from two positions. A system that is systemically arrestable with respect to its operator is, by the same architecture, arrestable with respect to anyone within reach of its effects — the halting capacity does not ask who is standing in the hazard zone. The operator-centred framing of this paper is therefore a methodological choice, not a limitation of scope: the operator is where the legal duties are textually anchored and where the temporal analysis is sharpest, and everything established there propagates outward without further argument.

5.2 The forthcoming Regulation as confirmation

Regulation (EU) 2023/1230 has appeared throughout this paper as evidence, and its role can now be stated precisely. The Regulation's requirement that, for control systems of machinery with fully or partially self-evolving behaviour, "it shall be possible at all times to correct the machinery or related product in order to maintain its inherent safety" (Annex III, section 1.2.1) is the first statutory provision to address the temporal problem in terms — and its language is instructive on every point this paper has argued. It speaks of *maintaining* inherent safety, confirming that continuous correctability is understood as already inherent in the safety such machines owe, not as a new layer added to it. It locates the requirement among the design requirements for control systems, confirming the migration of the duty described in Section 3. And it attaches to a category — self-evolving behaviour — confirming, by its very specificity, the need for the general criterion: the Regulation answers the problem for one class of machines, while the problem, as Section 3.4 showed, attaches to speed and arises wherever speed outruns reaction. On the reading proposed here, the Regulation is not the solution to the effectiveness gap but the legislator's first acknowledgment of it — and systemic arrestability is the principle of which the correction requirement is the first statutory instance.

5.3 The military question, noted and left open

The analysis has been confined to civil and dual-use machinery because that is where its anchors hold: machinery specially designed and constructed for military or police purposes is expressly excluded from the European machinery regime, and military systems likewise fall outside the LSPro's market-placement regime and OSHA's occupational scheme. The exclusion is jurisdictional, not logical. The temporal structure examined here — a hazardous capacity whose effects outrun any human response chain — does not change its nature at the boundary of civil law, and a person operating or accompanying such a system stands in the same factual position of objective exposure wherever the system is deployed. Whether, and through which instruments, an arrestability criterion has purchase beyond civil product-safety law is a question this paper deliberately does not develop. It is noted here for a single reason: the dual-use character of much contemporary machinery means the boundary will be crossed by the machines themselves, whatever the literature does.[13]

5.4 Open questions within the civil domain

Three questions internal to the argument deserve further work. First, quantification: the comparison of windows requires defensible figures for minimum credible human response time in defined configurations, and while the human-performance literature supplies orders of magnitude, conformity assessment will need consolidated reference values — a task for standardisation, guided by the legal criterion rather than substituting for it. Second, the residual class: there may be systems for which neither criterion can be satisfied — which can neither allow time nor be disabled within the window — and the legal status of such systems under existing duties (whether they can be considered safely designed at all) is the hardest downstream question this paper's framework generates. Third, burden and proof: if reliance on operator intervention must be justified by the window, the allocation of that justificatory burden in conformity assessment and in litigation merits treatment of its own.

6. Conclusion

Machine safety law has always protected the operator, and it has always done so against time: every stop function, every emergency device, every guard is a way of holding open an interval in which a human being can act. This paper has argued that the protection rests on a temporal assumption — a humanly bridgeable interval between hazard and harm — that is identifiable in the text and architecture of the European, Swiss, and American frameworks, and in the European texts is stated in terms; that a definable class of machines closes that interval while leaving the protective forms intact, producing not a regulatory gap but an effectiveness gap; and that the gap is closed by an interpretative criterion, *systemic arrestability*, under which a system's capacity to halt a hazardous action must not depend on human perception, evaluation, and intervention within the window of that action.

The criterion creates nothing. It consolidates a migration of duty that existing law already performs in fragments — in a design hierarchy, in voluntary standards, in category-specific rules — into a single proposition that manufacturers can apply, assessors can verify, and courts can enforce. Its

[13]Regulation (EU) 2023/1230, Art. 2(2)(l), excluding "machinery or related products specially designed and constructed for military or police purposes." An equivalent exclusion appears in Directive 2006/42/EC, Art. 1(2)(g).

instances already surround us, unremarked, in every interlock and every scram circuit. What has been missing is not the practice but the principle; and a protection that exists in practice but not as principle is exactly the kind of protection that fails silently, one formally compliant machine at a time. A machine that relies on human intervention as its safety mechanism, in conditions where human intervention is physically excluded, has not discharged the duty of safe design; it has only documented it. The operator whose hand can no longer reach the stop in time has not lost the right to a machine that stops. The law that promised the stop has only to be read as meaning it.

References

Legislation and regulations

Directive 2006/42/EC of the European Parliament and of the Council of 17 May 2006 on machinery, and amending Directive 95/16/EC (recast), OJ L 157, 9.6.2006, p. 24.

Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery and repealing Directive 2006/42/EC and Council Directive 73/361/EEC, OJ L 165, 29.6.2023.

Loi fédérale sur la sécurité des produits du 12 juin 2009 (LSPro) / Legge federale sulla sicurezza dei prodotti, RS 930.11.

Occupational Safety and Health Standards, 29 CFR § 1910.212 (General requirements for all machines), Subpart O – Machinery and Machine Guarding.

Technical Regulation of the Customs Union TR CU 010/2011, “On safety of machinery and equipment”, Decision No. 823 of 18 October 2011 (Eurasian Economic Union).

GB/T 15706 (ISO 12100, IDT); GB 5226.1 (IEC 60204-1, IDT); GB/T 16855.1 (ISO 13849-1, IDT) (People’s Republic of China, national adoptions).

Work Safety Ordinance (New Version), 5730-1970 (Israel).

Case law

Case T-337/13, *CSF Srl v European Commission*, judgment of the General Court (Third Chamber) of 15 July 2015.

Akron Brick & Block Co., 3 BNA OSHC 1876 (No. 4859, 1976) (OSHRC).

Rockwell International Corp., 9 BNA OSHC 1092 (No. 12470, 1980) (OSHRC).

Fabricated Metal Products, Inc., 18 BNA OSHC 1072 (No. 93-1853, 1997) (OSHRC).

Aerospace Testing Alliance, OSHRC Docket No. 16-1167, decision of 21 September 2020.

Official guidance

European Commission, *Guide to application of the Machinery Directive 2006/42/EC*, 2nd edition, June 2010.

European Commission, Machinery Working Group, document MD WG 2010-03 EN rev. 2, *Emergency Stop Devices (MD Annex I 1.2.4.3)*, approved 5–6 November 2014.

Literature

Green, M., "'How Long Does It Take to Stop?' Methodological Analysis of Driver Perception-Brake Times", *Transportation Human Factors*, vol. 2, no. 3 (2000), pp. 195–216.

Summala, H., "Brake Reaction Times and Driver Behavior Analysis", *Transportation Human Factors*, vol. 2, no. 3 (2000), pp. 217–226.