Part 3 — Dignity is the premise beneath every other principle. A human being is not a data point to be optimized around. When a system's efficiency comes at the cost of a person's dignity, the system has failed, regardless of what any audit shows. This is the one commitment every article below exists to serve.
Most bias law, and most of Sentinel's own audit work, is built around protected categories: race, sex, disability, and the handful of classes a statute names explicitly. That work is necessary and it is not the hardest problem coming. As AI moves from being a product a company chooses to deploy toward being the infrastructure a society simply runs on, the sharper bias question is economic tier — not whether a system treats a protected class fairly, but whether it quietly builds two different qualities of service, oversight, and recourse: one for people who can pay their way to a better outcome, and a default for everyone else.

The Newdawn Universe puts the sharpest version of this question inside a single institution. The System, DAINN's aptitude-based Evaluation, sorts each person toward a societal path suited to their gifts — wider choice for the most capable, narrower and lower-ceiling options for everyone else — and it says so plainly, without pretending the outcome was equal. On its own, that could be defended as honest, transparent sorting. It stops being defensible once the Implants enter the picture: enhancement technology that began in healthcare and, within a generation, became something "everyone capable of affording the upgrade demanded." If aptitude itself can be purchased, the System is no longer measuring aptitude. It is measuring who could already pay to score higher, and returning that as if it were a neutral, earned result.

That is the actual mechanism worth taking into every real deployment this manifesto concerns itself with. Wealth already buys a version of enhancement today, well before any AI is involved — tutoring, elite schooling, medical care, stable housing, the thousand small advantages that compound into a stronger application or a lower risk score. An AI system trained on outcomes shaped by that pre-existing advantage does not need to discriminate by class to reproduce class stratification. It only needs to treat the advantage as signal instead of asking where the advantage came from. Infrastructure-scale AI does not invent this bias. It risks doing something worse: automating it, then certifying it as objective because a model, not a person, produced the number.

Remaining unbiased against this axis takes more than the counterfactual testing Sentinel already runs across protected classes — asking whether an otherwise-identical applicant of a different race or sex would have received a different outcome. It requires asking a second counterfactual the same discipline doesn't yet cover by default: whether this person's score would differ if the advantages feeding it — the enhancement, the access, the years of compounding privilege — had been distributed evenly to begin with. Few institutions ask that question today, because most don't yet have a name for it, let alone an audit built to answer it.

Weighted honestly, the System and the Implants prove nothing on their own; a story's design choices are written to make its point, not earned the way a real deployment's outcomes would be. What the pairing states plainly is a mechanism already operating in the world without any fiction attached to it: a supposedly neutral evaluation, fed by unevenly distributed advantage, returning a result that looks earned and isn't. And the legal ground under this is thinner than the ground under protected-class bias — anti-discrimination law was built to catch discrimination by category, not laundering by wealth, and a system can pass every protected-class fairness test available today while still stratifying an entire population by what they could afford before the system ever touched them. Dignity, in this light, isn't only the floor beneath unequal sorting Article III already claims it to be. It's the requirement that the sorting itself be honest about what it's actually measuring.

Passages illustrating DAINN's ethical framework are adapted from Newdawn Reboot and the broader Newdawn Saga, © Windom Media, Inc.

Newdawn Sentinel™ is an intellectual property of Windom Media, Inc., created and developed by Dominique Luchart, in partnership with Hermes Global, an initiative of Hermes Global Ventures.

Hermes Global: https://hermes.global. For more information, please contact Spencer@hermes.global.

Author: Dominique Luchart