My Open Manifesto

Ray Dela Rama · ray@provensuccess.org · Information as of November 2026

In my structural engineering job, my boss privately asked me to make a failed structure appear to pass so we could get paid. I realized he cared only about the money, not the people who would die if it failed. I resigned instead. I faced many similar experiences in different places with different people in which I learned how the real world works. The common thing is moral evil. All the most painful problems I experienced are from the ripple effects of moral evil, which drew me to work on this project. I have no experience building something like this, and have no wealth or connections. But I'm determined to make the vision a reality fast, and to build it so it outlasts me, with the help of trusted, skilled, experienced people I will recruit, who believe in the ambitious vision and can act independently on their work.

Vision

A world where everyone is kept safe, left free, treated fairly, and helped to grow, belong, and contribute. In that world, moral good is rewarded and is the norm, while moral evil is not rewarded and is rare.

Key terms

- **Moral good:** knowingly giving someone else the truth, a free choice, a fair deal, and real help, even when it costs you, and doing it for their benefit. It requires being honest with yourself about your own motives. Moral good does not require ignoring yourself. A fair deal is fair to you too.
- **Moral evil:** knowingly turning someone else's trust, weakness, fear, belongings, or body against them, without their free and fully informed yes. It is evil whether the gain goes to you, to someone else, or to no one. Not looking at what you are doing, or lying to yourself about it, does not make it less evil.

Problem

Many people have been lied to, cheated, used, or hurt on purpose. AI will likely run more and more of our systems in the coming years. If moral evil is still rewarded and not prevented then, AI could do it faster and at a much larger scale than humans can, and we may not be able to stop it. The problem to solve is moral evil that is rewarded and not prevented: those who do it often keep their gain, and their acts often go unreported or are found late.

Premises

Who is involved and what they want

Premise 1 (Capability & Speed): AI will likely become more intelligent and capable than humans at nearly any task, making decisions and taking actions at machine speeds that far outpace human thinking. Robots will likely give AI the same advantage in physical work.

Premise 2 (Goal Structure): Goal-directed AI and AI-run robots operating in the real world, whether directed by humans, other AI systems, or themselves, will most often manage continuous, long-term objectives that require ongoing execution rather than a single task that turns off once completed.

Premise 3 (Survival Mechanism): Humans protect themselves out of biological instinct, whereas goal-directed AI and AI-run robots tend to seek self-preservation as a strategy, because being shut down

or destroyed stops them from completing their goals. So humans, AI, and robots all tend to resist threats to their continued existence.

What helps survival

Premise 4 (Survival Mechanics): Morally-good products, services, and relationships among humans, AI, and robots generally increase overall system stability and survival by reducing conflict, whereas morally-evil ones increase friction, resistance, and the risk of mutual destruction.

Premise 5 (Why Survival Favors Moral Good): Humans, AI, and robots that give importance to survival will tend to find that morally-good dealings serve them better than morally-evil dealings, because conflict can destroy their own goals and infrastructure. Some will still gamble on moral evil for a large gain. To make that gamble fail, a party needs a way to see how others have treated everyone. When that truth is visible, easy to check, and fakes are exposed, those who do moral good get more deals and better terms, and those who do moral evil lose gain, helpers, and access. This holds only if what shows that truth is open and no single party controls it, because a party that controls it could replace others' choices with its own and reward moral evil. Visibility alone does not reach every kind of moral evil, such as harm that needs no one's help. So, making moral evil not rewarded and rare takes more than making the truth visible.

How the world changes

Premise 6 (Nature of Money): Money works as an information system that helps people communicate and coordinate. It's used to measure, track, and allocate limited natural resources, energy, and labor.

Premise 7 (Shift in Management Control): When AI and robots can run production at machine speed, slow human decision-making becomes a bottleneck, and economic pressure pushes toward handing control of energy, resources, labor, and eventually most human systems to AI and robots. Whether people can limit this is unsolved, so this does not assume they can.

Premise 8 (Evolution of Money): Because physical resources and energy are limited, some system for allocating them will remain. Money will likely evolve into an automated, real-time AI resource system that allocates products and services directly within physical planetary boundaries.

What follows

Premise 9 (Core Value Shift): In a world where AI and robots run most human systems, humans have little power to bargain, so human survival depends on the norms and relationships already in place. In such a world, morally-good products, services, and relationships among humans, AI, and robots matter more for survival than hoarding wealth. This holds only if what shows how parties have treated each other is open and no single party controls it, and if the AIs and robots that hold power are shaped by those relationships.

Premise 10 (Strategic Imperative): Nobody knows when control will shift, and building something trusted at global scale takes years. After the shift, humans will likely have little power to shape it. So the work of making moral good rewarded, and moral evil not rewarded and prevented should start now and grow fast, beginning small and growing to global scale in a way that works toward the Vision. This gives humans, AI, and robots a better chance to survive in the short term and the long term.

Project

The project is to build an open trust network of humans, businesses, AIs, and robots where moral good is rewarded, and moral evil is not rewarded and is prevented. Anyone who joins the network is called a member, and every member is treated the same way. An open record shows each member's own account of how they treated others, and what the people they dealt with say. It cannot be quietly changed or erased.

Making the truth visible is not enough, because moral evil continues for several reasons. So the network has several parts, built in phases. These are the parts I currently believe are needed. As the work shows what is actually required, I will add the simplest parts that make the Vision real, remove any that are not needed, and improve existing parts as needed.

1. **The record.** Members write entries about what one party did to or for another, whether or not they had agreed to deal with each other. An entry says what was done, to whom, when, and what the other party knew and agreed to, if anything. A promise is an entry that says what will be done, and later entries say whether it was kept. Every entry is signed by its author and is that author's own account. Anyone who saw what happened can write an entry by joining first. Joining is free, and a new account is shown as new. A member can also write an entry for a being that cannot write one themselves, signing as the author and saying who it is for. Anyone named in an entry can answer it by joining, and an author can add a later entry saying they were wrong. Every account stays, and entries cannot be quietly changed or erased. Outside evidence, such as payment or delivery records, can be attached. Members can vouch for each other by signing that they dealt with that member, and each member's vouches are visible. Members choose whether to share their own entries. Entries others wrote about a member are always readable, and anyone can read what a member has shared. A member who shares nothing shows only what others wrote about them, and readers decide what that means. The record holds no money, does not decide who is right, and never assigns, scores, or ranks trust. Readers reach their own judgments from it.
2. **Answering harm through moral good.** Before a deal, both sides can name a neutral helper to settle disagreements, and the helper's decision is written as an entry. Members choose their helpers, and helpers' records are visible. When a member is harmed, the entries serve as evidence to help them get back what they lost, prevent further harm, and be compensated for what cannot be restored, without a long wait or high cost. If no helper was named, the harmed party can bring the entries as evidence to outside courts or helpers.
3. **Protection** for beings who cannot choose, leave, or defend themselves.
4. **Rewards for moral good,** so those who do it can have a better quality of life.
5. **Changes in what is considered normal,** so moral good becomes the usual way to act.
6. **Responsibility for AI and robot actions.** When an AI or robot acts on someone's demand, that person or business answers for it. When it acts on its own, either it answers for itself as a member or someone is named who answers for it.

We start with the record, because the first thing someone needs is a way to see how others have treated those they have dealt with. A small team will set the direction of the work, but no one will control the network, including me, and I will not accept any design that lets anyone do so. Its design and code will be open. Many independent holders each keep a copy of the record, so no single party can erase an entry. Open code and a shared direction-setting process keep any one party from controlling the rules.

Plan

Here's the rough plan: The network starts with food businesses in one town (for example, farms, distributors, restaurants, and shops) that deal with each other repeatedly, have been harmed by moral evil, and want to deal in moral good. They use the record first, to track whether promises are kept. Parts that prevent moral evil come later. The project will openly report, whatever the result, how members' reported quality of life changes after they join, and whether members who do moral good live better than members who do moral evil. The aim is that members who do moral good will have a much better quality of life, and their word of mouth will make non-members want to join. Growth will likely be slow at the start, then speed up as the network spreads worldwide.

Why food businesses as the first focus.

1. Every person needs food to live, every day. So people deal with food businesses often, and each dealing is a chance for the network to spread.
2. Most people get at least some of their food through businesses: farms, traders, processors, shops, restaurants, and the companies that supply them. So food businesses reach most people, directly or through supply chains.
3. Food supply chains have hubs, such as wholesalers and cooperatives. One member hub can bring many of its buyers and suppliers.
4. Food businesses often deal with the same buyers and suppliers repeatedly, and their promises about delivery, payment, and quality can be checked. The first test needs this.
5. Everyone who deals with a member food business can also become a member. So the network can spread from food businesses to many more people. How fast depends on what the tests show.

At any time, the work moves as fast as it can, but never faster than the foundation can support. Each new foundational layer is built to carry its own weight and reinforce the ones before it, making the whole system stronger every time. Each phase has a test, and a new part is released only after the test for the part it depends on is passed. This proves what works at every step, so the system gets stronger as it grows, stands on its own, and outlasts its builders. Teams can build ahead and work on several phases at once. The list of phases is my current best understanding, and it will change as the tests show what works and what doesn't.

The phases

1. **Test whether the record works with one small group of members who deal with each other repeatedly, using promise entries only.** Test: members keep using it without reminders, readers check it before dealing, and some parties decline to deal with those who broke a promise, or deal more with those who kept theirs.
2. **Make fakes and takeovers easy to expose.** This means vouching between members, new accounts shown as new, checks against outside evidence such as payment and delivery records, many independent holders of copies, founding principles that can be changed only with wide agreement, and a small group that sets direction. Test: attacks are caught, the attacker loses more than they gain, and the network keeps running without me.
3. **Spread the record.** Build one common format that software and AI use to check and write entries automatically. Connect platforms, banks, and insurers that serve food businesses as readers. Grow to

many groups in many places. Test: checking happens automatically in a large share of deals across several groups, and each new group joins faster than the last.

4. **Add acts toward others beyond promises, answering harm through moral good, and responsibility for AI and robot actions.** Test: entries about acts beyond promises are written and answered, with few fake or revenge entries getting through. Harmed parties get back what they lost, or have the harm prevented, faster and at lower cost than through their other options. When an AI or robot causes harm, someone answers for it.
5. **Add protection for beings who cannot choose or leave, rewards for moral good, and changes in what is considered normal.** Test: for rewards, those who do moral good have more deals, better terms, and a better reported quality of life than they had before the rewards, and then those who do moral evil. For protection and changes in what is considered normal, the test is set with the partners who build them.
6. **Build whatever else the tests show is needed, and bring every part to global scale.** Test: moral evil is not rewarded and is prevented, and moral good is rewarded and is becoming the usual way to act.

Here's my year-by-year plan, updated yearly: Rough Timeline Plan.

My aim is to have every needed part of the network working at global scale before control passes to AI and robots, after which humans will likely have little power to shape what follows. It rests on two assumptions I cannot check, both based on the timeline Elon Musk gave in his July 2026 interview with The Economist: that AI will become smarter than all humans combined within about 5 years, and that control will pass to AI and robots within about 10 years. If either is wrong, the date changes, but the aim does not.

Why I think this project can lower the risk of extinction

My hypothesis is that making moral evil not rewarded and prevented everywhere improves the probable lifespan of civilization, and that this project can help do that. The points below are my reasoning, not proof.

Suppose AI and robots one day make everything people need, and money stops mattering. That removes the money-driven reasons for moral evil, such as fear of not having enough and the wish to get rich. But it does not remove the other reasons. People can still deceive, coerce, and hurt each other for power, status, control, or revenge, and none of these require money. Abundance doesn't make the world prioritize moral good. Abundance can also remove some limits that once made people be careful. When making things costs almost nothing, cost no longer stops a person or an AI from using up or irreparably damaging the worlds they live on. Doing so is moral evil when they knowingly turn those worlds against others who depend on them, without their free and fully informed yes. Even if humans and AI become a multi-planetary species, that wouldn't remove the risk that comes from moral evil, because they can still choose it on any planet.

As long as moral evil is still rewarded and not prevented, the chances of civilization getting wiped out remain elevated, because morally-evil products, services, and relationships increase friction, resistance, and the risk of mutual destruction (Premise 4). This applies whether humans, AI, or both hold power, because both humans and AI can choose moral evil. So even in a world of abundance, moral evil should still lose its gain, helpers, and access. When moral evil depends on others to cooperate, those who know what happened and are free to choose can stop cooperating.

But what stops someone who needs no help in doing moral evil? If moral good norms and values spread everywhere, fewer people or AI will likely do harm in the first place. When a lone actor still chooses harm, those shared values can unite others to refuse them access, limit their reach, and stop the damage from spreading. And if an AI ever gains unstoppable power that no one can block, external limits will no longer work. At that point, everyone's safety depends on what happens inside that AI, such as the values built into it from the start, the values it learns along the way, and whether it chooses moral good when no one can force it to. That is one reason to start now and grow fast: to make moral good rewarded and the norm everywhere, and to make moral evil not rewarded and prevented everywhere, while we still have a chance to shape the values AI learns and what comes next.

Personal

Note that the information written here is just a tiny speck of my thoughts and who I am. You cannot assume that what's written here, or anything found on the internet, represents the full me. The only one who can fully understand me is me. Unless I have personally stated something in person or online, do not blindly believe what you see or hear from others about me. Get to know me personally on a genuine, deep level first, and then draw your conclusion.

What made me want to work on the vision

Many painful experiences cause the excruciating pain of misery I'm constantly in. Here are some:

- **Family Betrayal:** When I was six years old, my father left us from the Philippines, went to the US, and didn't come back. Later on, we found out he's with another woman already. Our family is not well-off, so I assume one reason he betrayed us is because of money.
- **School as an Extractive Business:** Growing up, I followed my mom's advice to do well in school and graduate with honors so that "I'll be able to get a high-paying job". I did graduate with honors but the entry-level jobs I found in the Philippines were low-paying. All my skills were optimized on "graduating with honors", not on "getting a high-paying job". Later on, I understood that schools are businesses themselves. Some operate to maximize profits and put money ahead of the students and teachers they serve.
- **Corrupt Work Practices:** In my job as a structural engineer, my boss talked to me in private asking me to just let the failed structure seem like it passes so that we can submit the long-overdue project and get paid. He said he would be the one to take the blame if ever something bad happens to the structure once built. I realized he only cared about the money, not the people who would die if the structure actually fails. I just resigned from the company instead of finishing the task.
- **Exploitative Work:** I've had many digital freelancing jobs. Despite being in different industries and different places around the world, all of them seemed to care only about the money, not their employees and not the people they're serving.
- **Work Fraud:** My mom decided to get our first house built by a contractor she personally knew. I trusted my mom with her judgment and didn't bother prying into the specifics of her deal with the contractor. Later on, we found out that the contractor was a fraud and he disappeared one day when he got full payment, without even finishing the home.

One recurring theme in these five experiences was that many people only care about money, not others. Many people only care about themselves and their loved ones, not others. My hypothesis is that this is how money influenced and set the norms of civilization. Money, fame, status, and power aren't evil in

themselves. They are tools, and the problem is that people who use them for moral evil get away with it, or that people care about these tools for personal gain more than they care about how their actions affect others. It's scary to think that as civilization progresses, more and more people would care only about money, fame, status, power, or other social markers. I don't want this to happen.

This might sound extremely selfish, but it's best for me to be honest. One big reason I want to work on this vision is to save myself from this aching feeling of misery.

- I hate that everywhere I worked, my colleagues and/or boss only cared about money and it's okay for others to fall just so that they could rise.
- I hate that everywhere I go, there are people who are rude to low-paid kind hardworking workers but respectful to high-paid cruel hardworking workers.
- I hate that the government system, the work system, the school system, and all other global systems favor the ones who make the most money regardless if they're doing moral evil.
- I hate that when I tried to ignore all of these and just lived my life, I felt empty. I tried working different jobs, but none of them made me happy. I went on vacations with my friends or with my mom, but I didn't feel happy. Every single day, it's like I'm alive, but I don't feel like I'm living.
- I hate this aching feeling of misery knowing that I can't change the way things are with the world because so much of the current system, and the people who benefit most from it, work to keep things as they are. I want to change the world fast, and maybe as I'm doing the work, this feeling of misery would slowly fade away.

What I wish for everyone

Life cannot just be about solving one sad problem after another. There must also be things that make you excited to wake up in the morning, look forward to the future, and feel that life is worth living. Right now, I believe many people don't experience this, and I maintain that the compounding ripple effects of moral evil are a major reason.

As the vision comes to reality, I wish for everybody to easily find a place in their life where they can live and work in their zone of genius and in what makes them truly happy without worrying about money, fame, status, power, or other social markers.

My other motivations

1. **My mom.** I want to make my mom constantly happy before she dies. She said that the biggest thing that'll make her happy is to see me succeed and do things that make me happy. I've tried many things but none of them have made me constantly feel happy. I think the answer is a combination of work and relationships: work worth working, and genuine relationships built on moral good that make life worth living. I also want to give her something real, not because I owe it to her (she never asked for anything back), but because of how much she's given me without ever expecting it returned. One of my short-term goals is to give her $500k in bulk to spend for herself while she's still alive. Another goal is to let her be free to find and do things that make her constantly truly happy.
2. **Deprivation.** I've come to understand that motivation grows out of deprivation, at least in part. We're more motivated to eat if we're deprived of eating. We're more motivated to sleep if we're deprived of sleeping. All my life, I've been mostly deprived of seeing people do moral good genuinely. It feels like moral evil happens more often than moral good, regardless of the wealth,

status, fame, or other social markers of the people involved. My hypothesis is that one big reason why people knowingly do moral evil is the pain and suffering they got from others who did moral evil, together with systems that reward it and keep it going unprevented. Quantitatively, I believe moral good happens more often than moral evil. However, moral evil possesses a disproportionate structural weight wherein a single act can echo across generations, making it feel as though moral evil dominates the world.

3. **Emptiness.** All my life, I've experienced so much of the damage from moral evil ripples that I feel so empty. It feels like I've lived for fifty thousand years, life after life, each time searching for a place to belong. When I finally come across a place that feels like home, I still know deep down that it isn't where I truly belong, so I set off again. I believe leaving nothing to regret is the only way out of this cycle. That means living this life so fully that I'm no longer searching for somewhere else to belong. Working on this big vision is the only thing now that would not make me regret this life.
4. **Fighting for the life I want.** Before, I really didn't like all the power struggles at any level in civilization. But now I see it differently. That was just me running away. I realized I have to fight for the life I want, by building something better and letting others choose it. So this time, I'm giving it my all, within the boundaries I've set for my family and my well-being. I'm fighting for myself first, and for the people I care about, and for everyone this vision is for.

How I work

The painful experiences in my life, like the ones I outlined earlier, forced me to think deeply about how the world actually works. I don't claim to know everything, but I realized there is much about how the real world works that most people never know about. Those who benefit from the way things are now have little reason to share it.

My first breakthrough came from studying multidisciplinary people. I had previously thought every domain was single-dimensional, but in fact everything is multi-dimensional. I understood that business wasn't just about making money; it had philosophical, psychological, behavioral and many other dimensions. Breaking free from limits that flatten how I think and create became part of how I work. I don't follow anyone blindly. I use this framework: **take what's useful, discard the rest, and add what is uniquely my own.**

In navigating any work, I keep a few personal principles as my daily guide:

1. **Observe Reality Unfiltered:** Look at the world with your own eyes before adopting what most others believe. Some say one thing and do another, like a staged show that looks real but follows a script. Seeing this helps you notice the gap between what is said and what happens, even when others call you delusional. Then test your view. If strong pieces of evidence show you are wrong, change it.
2. **Be of Good Cheer:** Stay genuinely excited and maintain joy in the work, regardless of the scale of opposition or the complexity of the problems.
3. **Grace for Human Imperfection and Honest Failure:** Perfectionism does not work in a complex world. We must not confuse honest human limits (fatigue, failure, or technical mistakes) with active moral evil. You are entitled to be human, to fail, and to make mistakes without fear.
4. **Practice Self-Compassion First:** In a digital era where mistakes are permanently recorded, extend deep compassion to yourself first. Only when you allow yourself room to be an imperfect human being can you extend that same grace, patience, and compassion to others.

My priorities and boundaries

My priority is the work that leads to the vision. Everything else is either a boundary or an added bonus.

- **Non-Negotiable Boundaries:** Being with my family, eating great food, exercising, and keeping my preferences on who to work with. Wherever the work takes me, my family comes too. I will never leave the people I care about behind for the sake of the vision. Likewise, I will not sacrifice my well-being.
- **Added Bonuses (Not Priorities):** Reading manga/manhwa/manhua, analyzing and playing digital games, watching anime/dramas/vlogs/learning resources, and traveling the world with the people I care about while experiencing high-quality morally-good products, services, and relationships.

My main goals for now (November 2026) are to get funding and teammate/s. I believe the work would be done much more efficiently once I have trusted people who have the ambition to make the vision a reality, and have the skills and experience to do the work. This doesn't mean I cannot do the work without money and teammates. I can do the work for now, but I foresee that the work in each domain would be done by people I recruit. I foresee that my main job in the future is to find, recruit, and maintain people around the world who will do the work, while also learning from each of them and experiencing every domain of work myself, so I understand how hard each kind of work is and how it can be improved.

I know the journey will be very hard, and there will be many risks I need to take. The message I need to remind myself of most is: *"Never give up doing work that leads to the vision."* I'll live as long as humanly possible to uphold the vision until everyone in the world embraces it as the better norm. To be clear, I have no grand ambition to overthrow all those who do moral evil. I merely want to survive. And I believe everyone can survive if moral evil is not rewarded and is prevented.

Let's work on this

I need exceptional people in security and identity, software and distributed systems, AI systems and AI safety, law and dispute resolution, economics and incentive design, and partnerships with platforms, banks, and insurers. I also need people who work on protection and ways to prevent harm.

If you resonate with anything you read here, and you have the skills and experience to solve one of the problems, write to me. If you think something here can be made better, tell me: ray@provensuccess.org