Emergency Response

Alex Applebee & L. N. Combe

2026

Civic Proximity Response and the Political Economy of Safety: A Unified Thesis on Community-First Emergency Infrastructure

Authors: Alex Applebee and L. N. Combe

Date: March 2026

Affiliation: OMXUS Research

Correspondence: research@omxus.com

Series: OMXUS Research Series, Paper No. 13

Version: Unified Thesis — consolidating the Evidence Synthesis, the Community Emergency Response and Political Economy of Policing analysis, the HREC Pilot Study Application, and the Children at Risk policy analysis

---

Author's Note

This paper exists because of a fall.

An elderly woman fell in her home. Her granddaughter lived sixty seconds away. The ambulance took fourteen minutes. Fourteen minutes is not a staffing problem. It is not a funding problem. It is an architectural problem. The system that determines whether your grandmother lives or dies was not designed to save your grandmother. It was designed to manage a fleet.

That granddaughter could have been there in sixty seconds. She was not alerted. She did not know. By the time she found out, the window had closed.

This paper is Goal 13 of the OMXUS project: *A $29 ring. Press it, your people come in 60 seconds.*

The idea is not original. Hatzolah has been doing it in Jewish communities since 1969 — volunteer medics who live in the neighbourhood, carry equipment in their cars, and arrive in under three minutes while the ambulance is still being dispatched. Volunteer surf lifesaving in Australia has been doing it on beaches since 1907 — ordinary people who train, who watch, who run toward danger because someone they can see is drowning. The model works. It has always worked. What

has never existed is the technology to extend it to every street, every home, every person who falls and has no one watching.

That is what we are building. A ring that costs twenty-nine dollars. You press it. Your people — not a call centre, not a dispatch algorithm, not a stranger in a uniform who has never been to your street — your people come. In sixty seconds. Because they are already there. Because they live next door. Because the system told them someone needs help, told them by name, told them how far away they are, and gave them a reason to move.

The evidence in this paper demonstrates that this is not optimism. It is mathematics. A Poisson spatial coverage model shows that in a suburb of 3,000 people per square kilometre, with only 4% adoption and a conservative willingness factor, you achieve 90% probability of a responder within 200 metres. That is a forty-second walk. The gap between that and the fourteen-minute ambulance is the gap between life and death for cardiac arrest, for choking, for severe bleeding, for a grandmother on the floor who cannot reach her phone.

The second half of this paper asks a harder question: if community-first response is faster, cheaper, and produces better outcomes, why does the current system exist? The answer is historical. Modern policing did not originate as community protection. It originated as economic enforcement — protecting mercantile interests in London, controlling enslaved populations in the American South, subjugating Indigenous peoples across the British Empire, breaking strikes for industrialists. The $237.7 billion the United States spends annually on policing is not a safety investment. It is the operating cost of a control apparatus that fails to solve the majority of crimes reported to it, kills over a thousand Americans per year, and incarcerates Black and Indigenous people at rates that would be called apartheid if they occurred in another country.

The alternatives documented in this paper — violence interrupter programs reducing shootings by 30-63%, mental health crisis teams handling thousands of calls with zero fatalities in 35 years, Housing First reducing homelessness by 88% — are not proposals. They are deployed, evaluated, and proven. The barrier to implementation is not evidentiary. It is political.

This paper is for the granddaughter. And for every person who has stood helpless while the system took fourteen minutes to do what love could do in sixty seconds.

— *A.A. & L.N.C.*

Table of Contents

Abstract

This thesis presents a unified evidence base for community-first emergency response and decentralised safety infrastructure. It integrates four bodies of work: (1) a mathematical evidence synthesis demonstrating that proximity-weighted community response systems can close the gap between emergency onset and critical intervention windows; (2) a comprehensive analysis of the Hatzolah volunteer emergency medical service model and comparable community-first responder programs worldwide, documenting response times of 2-4 minutes compared to 8-14 minutes for traditional EMS; (3) a political-economic history of modern policing tracing its origins not to community protection but to economic enforcement, racial control, and labour suppression, with evidence that contemporary policing in the United States ($237.7 billion, 2024) and Australia ($17.9 billion, 2024-25) produces poor outcomes relative to investment; and (4) a complete Human Research Ethics Committee application for a pilot study of a wearable-initiated community proximity emergency response network.

We present evidence from deployed community responder systems — GoodSAM, PulsePoint, Hatzolah, CAHOOTS, Cure Violence, community paramedicine — demonstrating that community-embedded, volunteer-driven, prevention-oriented safety infrastructure outperforms centralised, professionalised, enforcement-oriented systems across virtually every measurable dimension: response time, clinical outcomes, cost-effectiveness, community trust, and reduction of harm.

Alternative safety models documented herein include violence interrupter programs achieving 30-63% reductions in shootings, mental health crisis teams diverting thousands of calls from police at 2% of the police budget with zero fatalities in 35 years, restorative justice programs reducing recidivism by 27-32%, and Housing First reducing homelessness by 88% and police contacts by 40%.

The thesis proposes a wearable NFC ring system ($29 per unit) that enables community emergency response in under 60 seconds via BLE mesh networking, with a Poisson spatial coverage model demonstrating that 4% adoption in suburban areas achieves 90% coverage probability within 200 metres.

The HREC application (Version 3.0, February 2026) details a prospective observational cohort study with n=500 participants over 12 months, including comprehensive ethical frameworks, risk assessments, privacy architecture, domestic violence safety protocols, and community engagement plans.

Keywords: community emergency response, Hatzolah, volunteer EMS, policing history, police spending, violence interruption, CAHOOTS, restorative justice, decentralised safety, community health, NFC ring, mesh networking, proximity alert, spatial coverage model, pilot study, wearable technology

---

PART I: THE RESPONSE TIME GAP

Chapter 1: Introduction — The Structural Problem

The assumption that community safety is best delivered through centralised, professionalised, state-controlled institutions — standing police forces, government-operated ambulance services, hierarchically organised fire departments — is so deeply embedded in modern governance that it is rarely examined as an assumption at all. It is treated as a fact of social organisation, as natural and inevitable as gravity.

This thesis argues that this assumption is not only untested but likely wrong.

We advance two interconnected claims. First, that community-embedded, volunteer-driven emergency medical response — exemplified by the Hatzolah model operating across Jewish communities worldwide — consistently outperforms centralised EMS on the metrics that matter most: response time, clinical outcomes, community trust, and cost-effectiveness. Second, that modern policing did not originate as a community safety institution but as a mechanism for economic enforcement, racial control, and the suppression of labour — and that its contemporary form continues to reflect these origins, consuming vast public resources while producing poor safety outcomes and significant collateral harm, particularly to marginalised communities.

These two claims are not merely adjacent. They are structurally connected. Both challenge the same underlying ideology: that safety is something delivered to communities by the state, rather than something that emerges from communities themselves. The Hatzolah model demonstrates that when communities organise their own emergency response, they do it faster, cheaper, and with better outcomes. The history and economics of policing demonstrate that when the state monopolises the safety function, it tends to serve interests other than community wellbeing.

The implications are significant. If the evidence shows — as we argue it does — that community-first models outperform centralised ones for emergency medical response, and that policing is both historically rooted in oppression and economically inefficient at producing safety, then the rational policy response is not incremental reform of existing institutions but a fundamental reallocation of resources toward community-based safety infrastructure.

This is not a utopian argument. Volunteer fire services already protect the majority of land area in most developed nations. Community paramedic programs already reduce hospital readmissions by 40%. Violence interrupter programs already reduce shootings by 30-63%. Mental health crisis teams already handle tens of thousands of calls per year with virtually no need for police involvement. The infrastructure for community-first safety already exists in fragments. What is lacking is the

political will to assemble it into a coherent alternative.

1.1 Sequential Dispatch Latency

The modern emergency dispatch model (000/911/999) imposes irreducible sequential latency:

Step	Duration
Call processing	60-90 seconds
Dispatch decision	30-60 seconds
Travel time	4+ minutes (NFPA 1710)
On-scene setup	Variable
Total	**7-14 minutes (urban), 20-45 minutes (rural)**

> **Remark**: This latency is architectural, not a staffing problem. Each step consumes time regardless of efficiency. Adding more ambulances to a centralised dispatch model reduces average travel time but cannot eliminate the call-processing and dispatch-decision delays. The system is sequential; the emergency is immediate.

1.2 Critical Intervention Windows

Emergency Type	Window	Survival Within	Survival Beyond	Source
Cardiac arrest	4-6 min	40-50% with CPR	<10% without	AHA 2020
Choking	4-10 min	>90% with intervention	Rapidly fatal	ERC 2021
Severe haemorrhage	5-15 min	High with pressure	~33% mortality	Eastridge 2012
Anaphylaxis	5-30 min	>95% with epi	~1% fatal	Simons 2011
Stroke	Minutes	1.9M neurons/min lost	Cumulative loss	Saver 2006
Violent assault	Seconds	Interrupts harm	Harm continues	—
Paediatric cardiac arrest	Minutes	5% ROSC loss/min delay	Rapidly fatal	Frazier et al. 2021

The Gap: For time-critical emergencies, the intervention window is often shorter than the minimum possible centralised dispatch time. This is not a funding problem. It is a design problem. The architecture of centralised dispatch cannot close this gap regardless of how much money is invested. The gap can only be closed by having someone already there.

Chapter 2: Literature Review

2.1 Community Health and Volunteerism in Emergency Services

The relationship between community-embedded health services and improved health outcomes has been documented across multiple disciplines. Marmot's (2005) work on social determinants of health established that community cohesion, social trust, and access to local health resources are

stronger predictors of population health than the availability of acute medical services. Putnam's (2000) research on social capital demonstrated that communities with dense networks of voluntary association — including volunteer emergency services — exhibit lower mortality rates, better mental health outcomes, and greater resilience to health shocks.

The volunteer emergency services literature is substantial. Perkins (1990) documented that volunteer fire departments, which constitute approximately 70% of fire services in the United States and over 90% in countries such as Germany, Austria, and Portugal (International Association of Fire and Rescue Services, 2023), provide effective emergency response at a fraction of the cost of professional services. Simpson (2008) examined the Hatzolah model specifically, finding that community-embedded volunteer medical responders achieved significantly faster response times than traditional EMS in comparable service areas.

More recently, the community paramedicine literature has demonstrated that extending paramedic roles into preventive and follow-up care produces significant reductions in emergency department visits and hospital readmissions (Patterson et al., 2023; Chan et al., 2019). The PulsePoint and similar citizen responder applications have been shown to increase bystander CPR rates from approximately 43% to 57%, with associated improvements in cardiac arrest survival (Brooks et al., 2022).

2.2 Policing Studies

The critical policing literature spans several distinct but converging traditions. Historical scholarship by Hadden (2001), Reichel (1988), and Turner, Giacopassi, and Vandiver (2006) has documented the direct lineage between Southern slave patrols — first established in South Carolina in 1704 — and post-Civil War police departments in the American South. Vitale (2017), in *The End of Policing*, synthesised this historical evidence with contemporary data on police spending, clearance rates, and community harm to argue that policing is not a broken system in need of reform but a system functioning as designed — to protect property and enforce social hierarchies.

The economic analysis of policing has produced mixed results. Chalfin and McCrary (2018) found that a 10% increase in police hiring produces a 3-10% decrease in crime, though the mechanism appears to be deterrence (police presence) rather than enforcement (arrests). However, Legewie and Fagan (2019) found that these effects do not hold in predominantly Black communities, where additional policing increases arrests for minor offences without reducing serious crime. Sharkey, Torrats-Espinosa, and Takyar (2017) found that community-based nonprofit organisations focused on crime prevention and community development were at least as effective as additional police in reducing violent crime.

2.3 Abolitionist Scholarship

The abolitionist tradition in criminology — rooted in the work of Angela Davis (2003), Ruth Wilson Gilmore (2007), and Mariame Kaba (2021) — argues that the prison-industrial complex and policing apparatus cannot be reformed because their fundamental purpose is not community safety but social control. This tradition has generated significant empirical research on alternatives to policing, including violence interruption programs (Butts et al., 2015), mental health crisis response (Shapiro et al., 2021), and restorative justice (Sherman & Strang, 2007).

While the abolitionist position remains politically contentious, its empirical claims — that policing produces significant harm, that alternatives exist and function effectively, and that investment in

social infrastructure reduces crime more efficiently than investment in enforcement — are increasingly well-supported by evidence from multiple disciplines.

2.4 Synthesis

What emerges from these literatures is a consistent finding: community-embedded, prevention-oriented, trust-based approaches to safety outperform centralised, enforcement-oriented, authority-based approaches across virtually every measurable dimension. The Hatzolah model provides the clearest demonstration in emergency medical response. The policing literature provides the clearest demonstration of the costs of the alternative approach. This thesis brings these two bodies of evidence together in a unified analytical framework.

Chapter 3: The Hatzolah Model

3.1 History and Structure

Hatzolah (Hebrew: "rescue") is a network of volunteer emergency medical service organisations operating primarily within Jewish communities worldwide. The first Hatzolah was established in the Williamsburg neighbourhood of Brooklyn, New York, in 1969, when Rabbi Hershel Weber organised community members to provide rapid first response to medical emergencies after observing that municipal ambulance response times in his densely populated neighbourhood were unacceptably slow (Hatzolah, 2024).

The model spread rapidly. By 2026, Hatzolah organisations operate in communities across the United States (New York, New Jersey, Florida, California, Maryland, Pennsylvania, and others), Israel (where United Hatzalah operates a nationwide service with over 8,000 volunteers), Australia (Melbourne and Sydney), the United Kingdom, South Africa, and several other countries.

The organisational structure is distinctive in several respects:

Volunteer workforce. Hatzolah responders are community members who volunteer their time. They are not paid professionals. They hold other jobs — as teachers, accountants, shopkeepers, tradespeople — and respond to emergencies as they occur, typically carrying medical equipment in their personal vehicles or on specially equipped motorcycles (ambucycles).

Community embedding. Responders live and work within the communities they serve. This produces two critical advantages: geographic proximity (the closest responder is often minutes or even seconds from any emergency within the service area) and social trust (patients know and trust their responders, reducing barriers to calling for help and improving patient cooperation).

Trained certification. Despite the volunteer model, Hatzolah members are trained and certified emergency medical technicians (EMTs) or, in many cases, paramedics. Training standards meet or exceed those of professional EMS agencies. Members maintain their certifications through regular continuing education and drills (Simpson, 2008).

Technology-assisted dispatch. Modern Hatzolah organisations employ sophisticated GPS-based dispatch systems. When a call comes in, the system identifies the closest available responders and alerts them simultaneously via mobile app or radio. United Hatzalah's dispatch system, modelled on ride-sharing algorithms, achieves nationwide coverage in Israel with average response times under three minutes (United Hatzalah, 2024).

Two-tiered response. Most Hatzolah organisations operate a two-tiered response: first responders (individual volunteers in personal vehicles or on ambucycles) arrive within minutes to begin treatment, while a fully equipped ambulance follows. This ensures that basic life support — airway management, CPR, defibrillation, bleeding control — begins at the earliest possible moment.

3.2 Response Times

The single most important metric in emergency medical response is time. For cardiac arrest, every minute of delay in initiating CPR reduces the probability of survival by 7-10% (Larsen et al., 1993; PMC, 2025). For stroke, each minute of delay in treatment results in the loss of approximately 1.9 million neurons (Saver, 2006). For traumatic bleeding, the "golden hour" concept — while debated in its precise parameters — reflects the well-established principle that faster intervention produces better outcomes across virtually all acute medical emergencies.

Hatzolah organisations consistently achieve response times that are dramatically faster than traditional EMS:

United Hatzalah (Israel): Average response time of less than 3 minutes nationwide, with 90-second averages in metropolitan areas. This compares to an average of 9 minutes for Magen David Adom (Israel's national ambulance service) ambulances (United Hatzalah, 2024).

Hatzolah Melbourne (Australia): Average response time of 4 minutes within their primary service area, compared to Ambulance Victoria's average response time of approximately 7-15 minutes depending on location and priority (Hatzolah Melbourne, 2024; Simpson, 2008).

Hatzolah New York: Published research documented median response times of 2-3 minutes, with Hatzolah arriving before the municipal ambulance service in 83% (29 of 35) of cardiac arrest callouts studied (Kadish et al., 2007).

General pattern: Across all documented Hatzolah operations, first responders typically arrive within 2-4 minutes, compared to 8-14 minutes for traditional EMS services in comparable areas.

The explanation for this differential is structural, not individual. Traditional EMS operates from centralised stations. When a call comes in, an ambulance must be dispatched from the nearest available station, navigate to the scene, and then provide care. Hatzolah responders are already distributed throughout the community. The closest responder may be a block away, at home, at work, or in a nearby store. The dispatch system identifies this proximity and the responder simply walks or drives the short distance to the patient.

This is not a marginal improvement. For cardiac arrest, the difference between a 3-minute and a 12-minute response is the difference between a patient who has a meaningful chance of survival with good neurological outcomes and one who almost certainly does not.

3.3 Clinical Outcomes

The clinical implications of faster response times are well-established in the emergency medicine literature:

Cardiac arrest survival. The global average survival rate for out-of-hospital cardiac arrest (OHCA) is approximately 8.8% to hospital discharge (Yan et al., 2020). In systems with very fast response times and high rates of bystander CPR, survival rates can exceed 20-25%. Since United Hatzalah's inception, the rate of cardiac arrest deaths in Israel has decreased by 50%, according

to the Israel Heart Society (United Hatzalah, 2024). While this figure reflects multiple factors (including improved public CPR training and defibrillator access), the rapid-response model is a major contributor.

Ambulance response time and survival. A study of 7,623 OHCA events in Norway found that each one-minute reduction in ambulance response time was associated with significantly improved outcomes, with the greatest marginal benefit occurring in the first 5 minutes (Holmen et al., 2021). This finding directly supports the Hatzolah model: shaving minutes off response time in the critical early window produces the largest survival gains.

Paediatric cardiac arrest. Every one-minute delay in EMS on-scene resuscitation after out-of-hospital paediatric cardiac arrest lowers the return of spontaneous circulation (ROSC) by 5% (Frazier et al., 2021). Children are particularly vulnerable to response time delays because their cardiac arrest mechanisms differ from adults and their physiological reserves are smaller.

Smartphone-activated volunteer responders. A 2025 study published in the *Medical Journal of Australia* found that smartphone-dispatched volunteer responders arriving before ambulance services were associated with improved survival to hospital discharge for OHCA patients (MJA, 2025). A randomized clinical trial published in PMC (2022) confirmed that smartphone dispatch of lay volunteers to retrieve AEDs and respond to cardiac arrests improved defibrillation rates and clinical outcomes.

Chapter 4: Community First Responder Models Worldwide

The Hatzolah model is not an isolated phenomenon. Community-first emergency response exists in multiple forms across the globe. What follows is an evidence review of the major deployed models and their outcomes.

4.1 Volunteer Surf Lifesaving (Australia)

Australia's volunteer surf lifesaving movement, established in 1907, is one of the oldest community-first emergency response systems in the world. Surf Life Saving Australia (SLSA) operates with over 180,000 members across 314 clubs, performing approximately 11,500 rescues per year. Volunteers patrol beaches, perform rescues, administer first aid, and operate emergency radio networks — all without salary.

The model demonstrates a critical principle: when the response resource is embedded in the environment where emergencies occur, response times collapse to seconds rather than minutes. A surf lifesaver standing on the beach reaches a drowning swimmer in under 60 seconds. An ambulance dispatched to the same beach takes 8-14 minutes.

The SLSA model also demonstrates the scalability of volunteer emergency response. With over 180,000 active members, Australia's surf lifesaving system is one of the largest volunteer emergency organisations in the world. It operates without conscription, without salary, and without coercion. People volunteer because the purpose is self-evident: someone is drowning; you can swim.

4.2 GoodSAM (United Kingdom, Australia, Global)

GoodSAM (Smartphone Activated Medics) alerts trained nearby volunteers when an emergency occurs. Smith et al. (2020) documented that the system nearly halves traditional ambulance response times for cardiac arrest events. The platform operates in the UK, Australia, New Zealand, and other countries, and has been integrated with some national ambulance services.

GoodSAM demonstrates that the critical variable is not professional certification but geographic proximity. A trained volunteer who is 200 metres away will always arrive before a professional paramedic who is 3 kilometres away. The question is not "who is more qualified?" but "who is closer?" For cardiac arrest, the person who arrives first and begins CPR is the person who determines whether the patient lives — regardless of their credentials.

4.3 PulsePoint (United States)

PulsePoint operates in over 4,500 communities across the United States. When a cardiac arrest is detected via 911, the PulsePoint app alerts CPR-trained citizens within walking distance and directs them to the nearest public AED. Brooks et al. (2016, 2022) found that PulsePoint increased bystander CPR rates and improved time to first defibrillation.

PulsePoint's data provides direct evidence for the willingness factor used in this paper's spatial coverage model: a significant fraction of people, when given a clear, personal, direct alert about a nearby emergency, will move toward the emergency. The bystander effect — the well-documented tendency for people to assume someone else will help — is substantially overcome when the alert is personal and specific ("You are 47 metres from someone in cardiac arrest. The nearest AED is at the pharmacy on the corner.").

4.4 Community Paramedicine (International)

Community paramedic programs extend the role of paramedics beyond emergency response into preventive care, chronic disease management, and post-discharge follow-up. Programs in the United States, Canada, Australia, and the United Kingdom have demonstrated:

- 40% reduction in hospital readmissions (Choi et al., 2023)
- Significant reductions in emergency department utilisation
- Improved management of chronic conditions (heart failure, COPD, diabetes)
- Cost savings for health systems through reduced acute care utilisation

The community paramedicine model inverts the traditional emergency-response hierarchy. Instead of waiting for people to call 000 when a chronic condition becomes acute, community paramedics visit patients at home, monitor their conditions, adjust medications, and connect them with social services. Prevention replaces reaction.

4.5 Volunteer Fire Services (Global)

Volunteer fire departments constitute approximately 70% of fire services in the United States and over 90% in Germany, Austria, Portugal, and several other European nations (International Association of Fire and Rescue Services, 2023). The global volunteer firefighting workforce numbers in the millions.

This is not a historical curiosity. It is the dominant model of fire emergency response on the planet. The majority of fire emergencies worldwide are responded to by volunteers, not professionals. These

volunteers are trained, equipped, and organised — and they deliver effective emergency response at a fraction of the cost of professional services.

The volunteer fire service model is direct evidence that community-based emergency response scales. It works in rural areas and small towns. It works in suburban communities. It works across cultures, legal systems, and economic conditions. The claim that volunteer emergency response is idealistic or unproven is contradicted by the largest body of evidence available: the global fire service.

4.6 CERT (United States)

The Community Emergency Response Team (CERT) program, administered by FEMA, trains civilians in basic disaster response: fire suppression, light search and rescue, first aid, and emergency organisation. Over 600,000 Americans have completed CERT training since the program's inception.

CERT demonstrates that non-professionals can provide effective emergency response with modest training. The program's existence is itself an admission by the federal government that professional emergency services cannot cover all emergencies and that community-based response is a necessary supplement.

4.7 First Responder Networks (Scandinavia)

Sweden, Norway, and Denmark have implemented smartphone-based first responder alert systems at national scale:

Sweden (SMS Lifesavers / Heartrunner): Ringh et al. (2015) published results in the *New England Journal of Medicine* showing that mobile-phone dispatch of laypersons for CPR in OHCA cases significantly increased bystander CPR rates. The system dispatches trained volunteers within a 500-metre radius of a cardiac arrest.

Norway: The evidence from Holmen et al. (2021) on the relationship between ambulance response time and OHCA survival has directly informed Norwegian policy on community first responder integration.

Denmark: The TrygFonden Foundation has funded nationwide deployment of AEDs and a citizen responder app, producing one of the highest rates of bystander defibrillation in the world.

The Scandinavian experience demonstrates that community first responder systems can be integrated with national health infrastructure at scale, producing measurable improvements in outcomes without replacing or undermining professional services.

4.8 Indigenous Safety Practices (Australia)

Aboriginal and Torres Strait Islander communities in Australia have practiced forms of community-first safety for tens of thousands of years. Traditional conflict resolution, mutual aid, kinship-based protection networks, and community governance structures represent the longest-running models of decentralised safety in human history.

The Disputes Centre Australia (2024) has documented traditional Indigenous conflict resolution methods that operate without police, courts, or incarceration. These methods — mediation by Elders, community conferences, restorative dialogue, kinship obligations — produce outcomes that

the colonial justice system has failed to replicate: genuine accountability, community healing, and the restoration of relationships damaged by harmful behaviour.

The over-policing of Indigenous communities (33% of prisoners from 3% of the population in Australia) is not evidence that Indigenous communities need more policing. It is evidence that policing as an institution is structurally incapable of producing safety in communities it was designed to control.

4.9 Community Health Workers (Global South)

In low- and middle-income countries, community health worker (CHW) programs have demonstrated that decentralised, community-embedded health response produces better outcomes than centralised systems — not as a compromise due to resource constraints, but as a structurally superior model for primary and emergency care.

Ethiopia's Health Extension Programme deployed over 38,000 health extension workers to rural communities, reducing child mortality by 24% and increasing contraceptive use from 15% to 35%. **Brazil's *Agentes Comunit??rios de Sa??de*** programme employs over 250,000 community health agents who have been credited with reducing infant mortality by 32% in programme areas. **India's ASHA workers** (Accredited Social Health Activists) number over 900,000 and serve as the frontline health response for rural populations.

These programmes demonstrate that the principles underlying the Hatzolah model — community embedding, proximity, trust, prevention orientation — are not culturally specific to Orthodox Jewish communities. They are universal principles of effective health and safety response that manifest wherever communities are given the resources and autonomy to organise their own care.

4.10 St John Ambulance Volunteer Networks (Australia, UK, Global)

St John Ambulance has operated volunteer first aid services since 1877. In Australia, St John provides volunteer ambulance services in Western Australia and the Northern Territory as the primary emergency medical service, alongside volunteer first aid at events and community programs across all states.

The St John model is particularly relevant because it operates within the formal emergency services framework — these are not informal community networks but registered, regulated, and government-contracted volunteer organisations delivering primary ambulance services to large geographic areas.

Chapter 5: Spatial Coverage Model — The Mathematics of Proximity

5.1 Pure Poisson Baseline

We model participant locations as a homogeneous Poisson point process with intensity rho (participants per square kilometre).

Proposition (Geometric Coverage): The probability that at least one participant is present within distance d of any point is:

```
P_present(d) = 1 - e^(-rho * pi * d^2)
```

This follows directly from the void probability of the Poisson process. The probability that no points fall within a circle of radius d is e^(-rho * pi * d^2); the complement gives the probability of at least one point.

5.2 Willingness Discount

Physical presence does not equal response. Not every person near an emergency will respond. We introduce a willingness factor w in (0,1]:

Definition (Effective Responder Density):

```
rho_eff = w * rho
P_response(d) = 1 - e^(-w * rho * pi * d^2)
```

5.3 Estimating Willingness

Source	Estimate
GoodSAM/PulsePoint acceptance (trained volunteers)	40-80%
Upper bound for general population	w <= 0.5
Spontaneous bystander CPR (unalerted)	10-40%
Lower bound with direct alert	w >= 0.1
Central estimates	**w = 0.15 (pessimistic) to w = 0.30 (moderate)**

The willingness factor is the primary empirical unknown that can only be resolved through pilot deployment. The HREC application in Part IV of this thesis is designed specifically to measure this parameter.

5.4 Density Requirements (200m Coverage)

Willingness (w)	90% Coverage	95% Coverage
0.10 (pessimistic)	183/km^2	238/km^2
0.15	122/km^2	159/km^2
0.20	92/km^2	119/km^2
0.30 (moderate)	61/km^2	79/km^2
0.50 (optimistic)	37/km^2	48/km^2

Achievability: Sydney CBD (20,000/km^2) needs only 1.2% adoption for 95% coverage at w=0.10. A suburban area of 3,000/km^2 needs 4.0% adoption at w=0.20 for 90% coverage. These are not aspirational targets. They are achievable adoption rates for a free, useful, low-burden system.

5.5 Response Time Estimates

Components: 1. Alert propagation: <5s (BLE <1s, WiFi Direct 1-3s, Yggdrasil 2-5s) 2. Decision time: 5-15s 3. Travel time: d/v where v approximately 1.5-2.5 m/s (walking speed)

Expected Nearest Responder Distance:

```
E[d] = 1 / (2 * sqrt(rho_eff))
```

Results (w = 0.20):

Environment	Pop. Density	rho_eff	E[d]	Est. Response	vs. EMS
Dense urban	20,000/km^2	4,000	~8m	**15-25s**	7-14 min
Urban	5,000/km^2	1,000	~16m	**20-35s**	7-14 min
Suburban	3,000/km^2	600	~20m	**25-40s**	10-20 min

The mathematical result is unambiguous. In any environment with reasonable population density and modest adoption, the expected community response time is measured in seconds. The expected professional EMS response time is measured in minutes. The ratio is not 2:1 or 3:1. It is 20:1 to 40:1.

5.6 Explicit Limitations of the Model

This synthesis cannot prove: - Exact willingness curves (requires pilot data) - False alarm fatigue rates - Coverage under heterogeneous adoption (clustering, time-of-day variation) - Optimal alert radius tuning - Behavioural response to repeated alerts over time

These gaps can only be resolved through controlled pilot deployment. The HREC application in Part IV is designed to address precisely these unknowns.

Chapter 6: System Design — The Ring

6.1 Proposed Architecture

1. **Proximity alerts**: Broadcast to nearest verified participants
2. **Personal address**: "You (NAME) are 47 metres from someone who needs help"
3. **Low-friction interface**: NFC ring ($29) with mesh network propagation
4. **60-second target**: Community response before professional dispatch
5. **Cascade fallback**: If no acknowledgment within 45 seconds, expand alert radius

6.2 Why a Wearable NFC Ring

- **Removes smartphone barrier.** Your phone is not always accessible. It may be in another room, in a bag, out of reach. A ring is always on your hand.
- **Single-touch activation.** Triple-tap on the ring. Gross motor only. No fine motor control required. This matters because the people who most need to trigger an alert — someone falling, someone being assaulted, someone in anaphylaxis — are under acute stress with degraded fine motor control.
- **Mesh network propagation.** BLE to WiFi Direct to Yggdrasil. No cellular dependency. No internet required. The alert propagates through nearby devices regardless of cellular coverage.
- **Direct personal alert overcomes bystander effect.** The bystander effect research (Darley & Latane, 1968; Fischer et al., 2011) demonstrates that diffusion of responsibility prevents people from acting in emergencies when others are present. The ring system overcomes this

by making the alert personal and specific: not "someone call an ambulance" but "You, Sarah, are 47 metres from someone who needs help. Two others have been notified."

- **$29 per unit.** The ring is passive NFC — no battery, no charging, waterproof, multiple sizes. At $29 ($15 at bulk procurement), cost is not a barrier to universal adoption.

6.3 System Operation

When a participant activates an alert:

1. Ring transmits BLE signal to paired smartphone
2. Smartphone constructs alert packet (timestamp, coarse location, emergency type if specified)
3. Alert propagates via BLE and Wi-Fi Direct to nearby participants' devices
4. Nearby participants receive audio/haptic notification with coarse distance and direction
5. Responders acknowledge ("going", "calling 000", or "observing")
6. Acknowledgments visible to alerter ("2 people are coming, ETA <2 min")
7. Alert automatically expires after 30 minutes or manual resolution

The system operates alongside existing emergency services at all times. It does not replace 000. It fills the gap.

6.4 Cryptographic Architecture

The system enforces privacy architecturally, not by policy:

Function	Algorithm	Purpose
Identity key (IK)	Ed25519	Long-lived device identity
Derived key (DK)	HKDF-SHA256	Epoch-specific key derivation
Session key (SK)	Rotated every 15 min	Prevents temporal linkability
Alert encryption	XChaCha20-Poly1305	Authenticated encryption of alert content
Message encoding	CBOR (RFC 8949)	Compact binary serialisation

Privacy guarantees: - No identity in broadcasts - Session key rotation every 15 minutes (prevents tracking) - Relay blindness (forwarding nodes cannot read alert content) - Rate limiting (3 SOS alerts per 24 hours per identity) - Escalation thresholds (<2 acknowledgments in 45 seconds triggers wider relay)

These are not policies that can be changed by an administrator. They are cryptographic constraints enforced by the protocol. The system *cannot* be used for surveillance because the data structures required for surveillance do not exist in the protocol.

PART II: THE POLITICAL ECONOMY OF POLICING

Chapter 7: Origins of Modern Policing

7.1 The London Metropolitan Police (1829)

The standard narrative of modern policing begins with Sir Robert Peel's establishment of the London Metropolitan Police in 1829. The Peelian principles — policing by consent, minimal use of force, crime prevention rather than punishment — are taught in every police academy and cited in every policing reform proposal. Peel is described as the "father of modern policing," and the Metropolitan Police Act is presented as a civilisational advance.

This narrative, while not entirely false, is importantly incomplete.

The Metropolitan Police was established in a specific historical context: the rapid urbanisation and industrialisation of early 19th-century London. The city's population had exploded, and with it had come concentrations of the urban poor whose presence alarmed the propertied classes. Peel's framing of policing as crime prevention should be understood in this context: "crime" was substantially defined as threats to property and commercial activity. The Metropolitan Police Act was championed not by the urban poor who were most vulnerable to violence, but by merchants, industrialists, and property owners who wanted their commercial interests protected (Emsley, 2009).

The Metropolitan Police's jurisdiction was explicitly designed to exclude the City of London — the financial district — which maintained its own separate police force (the City of London Police). This arrangement ensured that the financial elite retained direct control over the force protecting their most concentrated assets while the Metropolitan Police managed the potentially disruptive populations of the surrounding boroughs.

The Peelian principles themselves may be partially apocryphal. Recent scholarship by Lentz and Chaires (2007) and later historical analysis published in *Policing and Society* (2024) has questioned whether Peel actually authored the nine principles commonly attributed to him, noting that they first appeared in print decades after the Metropolitan Police was established and may represent a retrospective idealisation rather than a founding document.

7.2 Slave Patrols in the American South

The history of policing in the American South follows a different and more explicitly brutal trajectory. The first slave patrols were established in South Carolina in 1704, more than a century

before Peel's Metropolitan Police. They spread throughout the slaveholding colonies and, after independence, throughout the slaveholding states (Hadden, 2001; Reichel, 1988).

Slave patrols served explicit functions:

- Enforcing curfews on enslaved people
- Checking travellers for permission passes
- Catching enslaved people who assembled without permission
- Preventing organised resistance
- Returning escaped enslaved people to their enslavers
- Terrorising enslaved communities into compliance through violence and threat of violence

The NAACP's historical analysis states plainly: "The origins of modern day policing can be traced to the 'Slave Patrol' " (NAACP, 2024). The scholarly debate about whether modern American police directly descended from slave patrols or merely share common characteristics is less important than the structural observation: in the American South, the institution that would become professional policing was explicitly designed to maintain a racial hierarchy for economic purposes. After the Civil War, Southern police departments often carried over aspects of the patrols, including systematic surveillance, the enforcement of curfews (now applied through vagrancy laws and Black Codes), and racial criteria for officer recruitment (American Bar Association, 2021).

7.3 Colonial Policing

The British Empire's approach to policing its colonies provides further evidence that the purpose of organised law enforcement was control rather than community safety.

Colonial police forces were modelled not on Peel's Metropolitan Police (which at least paid lip service to community consent) but on the Royal Irish Constabulary, established in 1822 — an "outwardly militaristic" force that was "involved more in dealing with political protest than the prevention and detection of crime." When London needed to establish police forces in its colonies, it looked to the Irish model rather than the English one, producing more harsh and militaristic agencies explicitly designed to suppress local populations.

In Australia, colonial mounted police forces were instrumental in facilitating Indigenous dispossession. They "facilitated Indigenous people's subjugation to colonial law and aided the establishment of agricultural and early industrial capitalist economies through land acquisition and settlement" (Nettelbeck & Smandych, 2010). The massacres carried out by colonial police and paramilitary forces against Aboriginal Australians are among the darkest chapters in Australian history.

7.4 The Strike-Breaking Function

The role of police in suppressing organised labour further illuminates the economic function of policing.

In the United States, police, the National Guard, and the U.S. Army played integral roles in suppressing the Great Strike of 1877, the Homestead Strike of 1892, the Pullman Strike of 1894, the Lawrence Textile Strike of 1912, and dozens of other major labour actions. Private detective agencies — most notably the Pinkertons — were hired by industrialists to infiltrate unions, provoke violence, and break strikes, often working in close coordination with public police forces (Harvard Political Review, 2021).

By the middle of the 20th century, the American police had evolved into what labour historians describe as "an autocratic, militarized force whose primary role was to challenge organized labor through union-busting and strike-breaking." The Memorial Day Massacre of 1937, in which Chicago police shot and killed 10 striking steelworkers, is perhaps the most notorious example, but it was far from isolated.

This history is relevant because it demonstrates a pattern: policing has consistently been deployed to protect the interests of capital against the demands of labour. When workers organise for better wages, shorter hours, or safer conditions — all of which are community safety issues — police are not deployed on the side of community safety. They are deployed on the side of the employer.

The structural argument is not that individual police officers are consciously serving capital. It is that the institution of policing, as historically constituted, has functioned as an enforcement arm of economic power. This function is not incidental to policing; it is foundational.

Chapter 8: Modern Policing — Costs vs. Outcomes

8.1 Police Spending Trends

United States. Total government police expenditures reached $237.7 billion in 2024, according to the U.S. Bureau of Economic Analysis (FRED, 2024). State and local police spending alone was $178.9 billion. The United States spends more than twice as much on police as on federal education.

Individual city budgets are staggering: New York City's NYPD budget exceeds $5.75 billion, Los Angeles spends over $3.2 billion, and Chicago over $1.9 billion. These are larger than the total government budgets of many nations.

Australia. Total real recurrent expenditure on police services was $17.9 billion in 2024-25, with an average annual growth rate of 2.0% over the preceding five years (Productivity Commission, 2026). New South Wales alone invests over $5.3 billion in police and community safety. When corrective services are included, the total criminal justice spend approaches $19 billion.

Trend. Police budgets have grown consistently in both countries over decades, even as crime rates have generally declined. The "defund the police" movement of 2020 produced almost no sustained reductions in police budgets. In most major U.S. cities, police spending in 2024-25 exceeded pre-2020 levels (Police Funding Database, 2024).

8.2 Crime Clearance Rates

Given the scale of investment, what do police actually achieve?

United States (2024, FBI UCR data): - Violent crime overall: 43.8% clearance - Murder and non-negligent manslaughter: 61.4% - Aggravated assault: 49.1% - Robbery: 30.4% - Rape: 27.0% - Property crime overall: 15.9% - Burglary: 15.2% - Larceny-theft: 17.3% - Motor vehicle theft: 9.2%

More than half of all violent crimes reported to police are never solved. For property crime — which constitutes the vast majority of all crime — more than 84% of cases are never cleared. The

national homicide clearance rate has fallen from 72% in 1980 to 61% in 2024. For rape, clearance has fallen from 49% to 27% over the same period.

To state this plainly: the United States spends $237.7 billion per year on policing, and police fail to solve the majority of crimes reported to them.

Context. Clearance rates measure only reported crimes. The Bureau of Justice Statistics estimates that more than half of all violent crimes and approximately two-thirds of all property crimes are never reported to police at all. When unreported crimes are included, the effective "solution rate" for all crimes is in the single digits.

The detection problem compounds these failures. Even the cases that police do investigate rely heavily on suspect interviews, witness credibility assessments, and confession evaluation — mechanisms that Paper 11 in this series (*Signal Inversion*) demonstrates are fundamentally broken. Human deception detection accuracy is 54% (Bond & DePaulo, 2006; N=24,483), barely better than a coin flip. Of the 23 behavioural cues people use to assess credibility, 91.3% are inverted: what investigators believe indicates deception actually indicates truth-telling ($p<0.0001$). False confessions account for 12-30% of documented exonerations. The $237.7 billion policing apparatus not only fails to solve most crimes, but cannot reliably determine truth in the cases it does pursue.

8.3 The Deterrence Myth

The standard justification for police spending is deterrence. The evidence is more nuanced than proponents suggest.

The most favourable research finds that a 10% increase in police hiring produces a 3-10% decrease in crime (Chalfin & McCrary, 2018; Evans & Owens, 2007). Hot-spot policing can produce localised reductions of up to 34% (Weisburd et al., 2016). These findings support the proposition that police presence has *some* deterrent effect.

However:

Diminishing returns. The United States already has approximately 900,000 sworn law enforcement officers. The marginal deterrent value of additional officers, at current saturation levels, is likely small and declining.

Racial asymmetry. Cities with the largest Black populations do not see the same policing benefits. Additional police officers in these cities do not appear to lower the homicide rate, while they do result in more arrests of Black people for low-level crimes (Legewie & Fagan, 2019).

Mechanism. The deterrence effect appears to operate through police *presence* (being visible) rather than police *action* (making arrests). This raises a fundamental question: if deterrence comes from visible presence rather than enforcement capability, is an armed police force the most efficient way to produce visible presence?

8.4 Police Violence as a Public Health Issue

Scale. Police in the United States kill more than 1,000 people per year, with 2024 marking the deadliest year since data collection began, with 1,365 police killings (Mapping Police Violence, 2025). On average, law enforcement kills someone every 6.44 hours.

Racial disparity. Black people are 2.9 times more likely than white people to be killed by police. Native Hawaiian and Pacific Islanders face the highest racial disparity at 7.6 times more likely

(Mapping Police Violence, 2025). The Lancet's comprehensive analysis of fatal police violence from 1980-2019 found that these disparities have persisted across decades and are not explained by differences in crime rates (GBD Police Violence Collaborators, 2021).

Accountability gap. Since 2005, only 98 non-federal law enforcement officers have been arrested in connection with fatal on-duty shootings. Only 35 have been convicted, often of lesser charges. Only three murder convictions have been sustained (Stinson, 2020).

Mental health impact. Police killings of unarmed Black Americans are responsible for more than 50 million additional days of poor mental health per year among Black Americans — a mental health burden comparable to diabetes (Bor et al., 2018).

8.5 Over-Policing of Marginalised Communities

Indigenous Australians. Despite comprising only 3% of Australia's population, Indigenous Australians account for 33% of the prison population — an all-time high as of 2023. The Indigenous imprisonment rate is 2,630 per 100,000 adult population, approximately 17 times the non-Indigenous rate (ABS, 2025). In the Northern Territory and Western Australia, First Nations children are incarcerated at 25 and 21 times the rate of non-First Nations children, respectively.

Black Americans. Black Americans represent 13% of the U.S. population but 37% of the prison and jail population (Sentencing Project, 2024). Black Americans are 3.6 times as likely as whites to be arrested for marijuana possession despite comparable usage rates.

These disparities are not evidence that policing is "broken." They are evidence that policing is doing what it has always done: surveilling and controlling populations that those in power perceive as threatening to social and economic order.

Chapter 9: Alternative Safety Models That Work

9.1 Violence Interrupter Programs

The Cure Violence model treats gun violence as an infectious disease and deploys credible messengers — community members with lived experience of violence who have transformed their own lives — to detect and interrupt conflicts before they escalate.

- **Chicago:** 52% reduction in killings (Skogan et al., 2009)
- **New York City:** 63% reduction in shootings (Picard-Fritsche & Cerniglia, 2013)
- **Baltimore:** Up to 56% reduction in killings and 44% reduction in shootings (Webster et al., 2012)
- **Philadelphia:** 30% reduction in shootings (Roman et al., 2018)
- **Cali, Colombia:** 74% reduction in killings (Cure Violence Global, 2024)

For every dollar invested in Cure Violence, cities save up to $18 in reduced medical and criminal justice system costs. This 18:1 return dwarfs the return on police spending by any metric.

9.2 Mental Health Crisis Teams

CAHOOTS (Crisis Assistance Helping Out On The Streets) in Eugene, Oregon, has operated since 1989. Key data:

- Responds to approximately 17% of calls that would otherwise go to police
- In 2019, handled 18,583 calls (60+ per day)
- Police backup needed for less than 1% of calls
- Costs approximately $2 million annually — 2% of the police budget
- Saves an estimated $8.5 million per year
- **Zero people killed by CAHOOTS responders in over 35 years of operation**

The CAHOOTS model has been replicated in dozens of cities. Denver's STAR program, Portland's Portland Street Response, and similar programs have produced comparable results.

9.3 Restorative Justice Programs

Meta-analytic evidence:

- 27% reduction in recidivism compared to traditional punitive measures (Latimer, Dowden & Muise, 2005)
- Victim-offender mediation reduces reoffending by 32% (Umbreit et al., 2001)
- 2025 meta-analysis of 46 studies with nearly 23,000 participants confirmed statistically significant reductions in general recidivism (Fulham et al., 2025)
- Significantly higher victim satisfaction than traditional court processes

9.4 Community Mediation

- Community mediation resolves 60-80% of cases that enter the process
- Cost per case: $100-500 vs. thousands for police/court
- Higher participant satisfaction than police or court outcomes

9.5 The Prevention Dividend

Housing First: 88% reduction in homelessness, 40% reduction in police contacts, 40% reduction in arrests (NLIHC, 2024). Every $10 invested returns $9.60-$21.72 in savings.

Early childhood programs: 19:1 return on investment through reduced crime and higher earnings (Heckman et al., 2010).

Community organisations: Every 10 additional community-based nonprofits per 100,000 people is associated with a 12% reduction in the murder rate and 10% reduction in violent crime (Sharkey et al., 2017). This effect is comparable to or greater than adding additional police officers.

Substance use treatment: 7 times more cost-effective than enforcement-based approaches (RAND Corporation, 1994).

PART III: SYNTHESIS

Chapter 10: Toward Community-First Safety Infrastructure

10.1 The Mesh Model

The evidence reviewed in this thesis points toward a model of community safety we term the "mesh model": distributed, community-embedded, technology-assisted, and oriented toward prevention rather than enforcement.

Distribution. Safety resources are distributed throughout the community rather than concentrated in centralised facilities.

Community embedding. Responders are members of the communities they serve. They know the people, the geography, the social dynamics.

Technology assistance. GPS-enabled dispatch, mobile health applications, citizen alert systems, telehealth platforms, and mesh networking support community responders with the information and coordination tools they need.

Prevention orientation. The primary goal is not to respond to emergencies after they occur but to prevent them from occurring.

Multi-tier response. Different situations require different responses. A cardiac arrest requires trained medical responders. A mental health crisis requires crisis counsellors. A neighbour dispute requires a mediator. The mesh model matches the response to the situation rather than sending the same response (armed police officers) to every call.

10.2 Reallocation of Resources

Consider a hypothetical reallocation of just 25% of US police budgets ($59.4 billion) toward:

Investment	Amount	Evidence Base
Community emergency medical response (Hatzolah model)	$5B	2-4 min response vs 8-14 min
Mental health crisis teams (CAHOOTS model)	$10B	0 fatalities in 35 years
Violence interrupter programs (Cure Violence)	$5B	30-63% reduction in shootings

Investment	Amount	Evidence Base
Community paramedic programs	$5B	40% fewer readmissions
Housing First	$15B	88% reduction in homelessness
Community mediation and restorative justice	$5B	27-32% reduction in recidivism
Early childhood and education programs	$10B	19:1 ROI
Substance use treatment expansion	$4.4B	7x more cost-effective than enforcement

This leaves 75% of police budgets intact for the functions where police are genuinely needed: responding to violent crimes in progress, investigating serious offences, and executing warrants.

10.3 Training and Certification

A community-first safety infrastructure requires robust training. The Hatzolah model demonstrates that volunteers can be trained to EMT and paramedic standards while maintaining other employment. The Cure Violence model demonstrates that effective violence prevention does not require law enforcement credentials. CERT demonstrates that non-professionals can provide effective emergency response with modest training.

The infrastructure for training already exists. What is lacking is the funding and political will to scale it.

Chapter 11: Children at Risk — The Policy Failures

> "We tolerate monstrosity when it is slow and profitable."

11.1 The Question Nobody Asks

Why is the cheapest dinner in Australia a Happy Meal?

Why does ice cream — marketed to children — contain ingredients linked to cancer, ADHD, and metabolic disease?

Why is bacon in the same cancer classification as cigarettes, but cigarettes have warning labels and bacon has recipe suggestions?

Why is the leading cause of death in Australia (heart disease) directly linked to diet, and yet we subsidise the industries that cause it?

Because we do not vote on it. Politicians funded by food lobbies vote on it. And they vote for profit, not health.

11.2 What Kills Australians

Cause of Death	Annual Deaths (approx)	Preventable?
Heart disease	18,000+	Largely yes
Dementia/Alzheimer's	15,000+	Partially
Stroke	8,500+	Largely yes
Lung cancer	8,500+	Largely yes
Diabetes complications	4,800+	Largely yes
Colorectal cancer	5,300+	Partially

Source: Australian Bureau of Statistics, Causes of Death 2022

Healthcare spending: $220+ billion/year (2022-23). Preventative health: ~2% of health budget. For every $1.40 invested in prevention, we save $13 in treatment costs (RACGP). But we do not invest in prevention. Because sick people are profitable.

11.3 Bacon and Cigarettes: Same Classification, Different Treatment

The World Health Organization's IARC classifies processed meat as **Group 1 carcinogen** — the same category as tobacco smoking, asbestos, and plutonium.

> "Processed meat was classified as carcinogenic to humans (Group 1), based on sufficient evidence in humans that the consumption of processed meat causes colorectal cancer."
>
> — IARC/WHO, 2015

Cigarettes: Graphic health warnings covering 75% of packaging. Advertising banned. Plain packaging mandated. Smoking rates dropped from 24% (1991) to 11% (2022).

Processed meat: No warning labels. Advertised freely, including to children. Promoted as "part of a balanced diet." Consumption remains high.

The difference is not scientific. It is political. The food industry lobby is more powerful than the tobacco lobby was before decades of litigation broke its grip.

11.4 McDonald's: 2.5 Million Happy Meals Per Day

A Happy Meal is often the cheapest dinner option for a family. That is not an accident. That is policy.

Agricultural subsidies flow to corn, soy, and wheat (processed food inputs). Vegetables receive almost none. The health costs of processed food are externalised — paid by individuals and the healthcare system, not by the companies that profit. McDonald's marketing budget: ~$2 billion/year globally. Broccoli's marketing budget: $0.

Result: A burger costs less than a salad. Feeding your kids poison is cheaper than feeding them health.

11.5 The Regulatory Capture

Food Standards Australia New Zealand (FSANZ) sets food safety standards. Their board and advisory committees include representatives from the food industry, the agricultural industry, and

the manufacturing industry. The people profiting from the current system decide what is "safe" in the current system.

Much nutritional research is funded by food companies. Studies funded by industry are **7-8 times more likely** to show favourable results for the sponsor (Chartres et al., 2016).

11.6 The Connection to Emergency Response

This is not separate from the emergency response thesis. It is why the emergency response thesis matters.

When a grandmother falls and the ambulance takes fourteen minutes, that fall may have been caused by osteoporosis accelerated by a lifetime of dietary calcium displaced by phosphoric acid in soft drinks subsidised by agricultural policy written by food industry lobbyists elected by a system that does not allow the grandmother to vote on food policy.

The emergency response gap is the acute symptom. The chronic disease is a governance system that does not allow the people harmed by bad policy to change it.

The ring solves the acute problem: someone comes in sixty seconds. Direct democracy solves the chronic problem: the people who eat the food decide what goes in it.

Chapter 12: Discussion and Limitations

12.1 Hatzolah Generalisability

The Hatzolah model operates within specific community contexts — primarily Orthodox Jewish communities with high social cohesion, strong norms of mutual aid, and dense residential patterns. Whether the model can be replicated in communities with less social cohesion or more dispersed populations is an open question.

However, the success of PulsePoint (4,500+ communities), community paramedic programs (international), volunteer fire services (70% of US fire departments), SLSA (180,000+ members in Australia), and community health workers (millions globally) demonstrates that the underlying principles — community embedding, volunteer motivation, rapid response — are transferable even if the specific organisational form is not.

12.2 Policing Complexity

The relationship between policing and safety is more complex than a simple "policing doesn't work" narrative suggests. Police do solve crimes, do deter some criminal behaviour, and do provide a necessary response to violent situations in progress. The argument is not that policing produces no value but that its value is far lower than its cost suggests, and that alternative investments would produce greater safety per dollar.

12.3 Selection Effects

Violence interrupter programs, restorative justice, and other alternatives may benefit from selection effects: participants who volunteer for these programs may be more amenable to change than those

who do not. Randomised controlled trials are difficult to conduct in criminal justice contexts, and observational studies may overestimate program effects.

12.4 Implementation Challenges

The gap between evidence-based programs in controlled settings and real-world implementation at scale is significant. CAHOOTS has succeeded in Eugene, Oregon (population ~175,000); whether it can scale to cities of millions is uncertain. Cure Violence has produced dramatic results in some neighbourhoods but more modest results in others.

12.5 The Willingness Factor

The single most important empirical unknown in the spatial coverage model is the willingness factor w. The existing evidence (GoodSAM acceptance rates, PulsePoint response data, bystander CPR rates) provides bounds but not precise estimates for the specific system proposed here. The pilot study in Part IV is designed to measure this parameter directly.

12.6 Political Feasibility

The political obstacles to community-first safety infrastructure are substantial. Police unions are among the most politically powerful organisations in most jurisdictions. Fear-based politics dominates public discourse on safety. Existing institutions resist reallocation. These obstacles are real but they are political, not evidentiary. The evidence for community-first models is strong. The barrier is will, not knowledge.

Chapter 13: Conclusion

This thesis has presented evidence for two interconnected propositions.

First, community-embedded emergency response — exemplified by the Hatzolah volunteer medical service — is faster, cheaper, and produces better clinical outcomes than centralised, professionalised EMS. Hatzolah's 2-4 minute response times, compared to 8-14 minutes for traditional ambulance services, translate directly into lives saved, particularly for time-critical conditions such as cardiac arrest, where every minute of delay reduces survival by 7-10%. The Hatzolah model is not an anomaly; it is part of a broader pattern in which community-based emergency services — including volunteer fire departments, citizen CPR alert systems, community paramedic programs, surf lifesaving, and community health workers — outperform their centralised counterparts.

Second, modern policing originated not as community protection but as economic enforcement. The London Metropolitan Police was established to protect mercantile interests. American policing in the South descended from slave patrols. Colonial policing was designed to subjugate Indigenous populations and secure economic extraction. The strike-breaking function positioned police as capital's enforcement arm. Contemporary policing continues to reflect these origins: consuming $237.7 billion annually in the United States and $17.9 billion in Australia while failing to solve the majority of crimes, killing over 1,000 Americans per year, and disproportionately surveilling, arresting, and incarcerating Indigenous Australians (33% of prisoners from 3% of population) and Black Americans (37% of prisoners from 13% of population).

The alternatives are not hypothetical. Violence interrupter programs reduce shootings by 30-63%. Mental health crisis teams handle thousands of calls annually with virtually no need for police backup. Restorative justice programs reduce recidivism by up to 27-32%. Housing First reduces homelessness by 88% and police contacts by 40%. Community paramedic programs reduce hospital readmissions by 40%. These programs work. They are cost-effective. They are replicable. They produce better outcomes for communities than policing does.

The question is not whether community-first safety infrastructure can work. The evidence shows that it already does, wherever it has been tried. The question is whether societies will choose to invest in what works or continue to invest in what does not.

The Hatzolah volunteer, who leaves dinner to respond to a neighbour's cardiac arrest and arrives in three minutes, is not performing an act of charity. She is demonstrating an alternative model of social organisation — one in which safety is something communities produce together rather than something the state imposes upon them. That model is faster. It is cheaper. It saves more lives. And it does not require anyone to carry a gun.

The ring costs $29. Press it. Your people come in sixty seconds. The mathematics works. The evidence is clear. The only remaining question is whether we build it.

PART IV: PILOT STUDY

Chapter 14: HREC Application — Civic Proximity Response Pilot

14.1 Project Summary

Title: Civic Proximity Response: A Pilot Study of Wearable-Initiated Community Emergency Networks in Suburban Australia

HREC Protocol Number: [TO BE ASSIGNED]

Version: 3.0 — February 2026

Principal Investigator: [NAME], [INSTITUTION]

Contact: research@omxus.com

Plain Language Summary

We want to test whether a simple wearable device (an NFC ring) can help people in emergencies get help faster from nearby community members, while ambulance and police services continue operating as normal.

Participants in one suburban community will receive an NFC ring and a smartphone app. If they experience or witness an emergency, they can tap the ring to send an alert to other participants nearby. Nearby participants who receive the alert can choose to respond by going to help, or not. The study measures how quickly someone arrives compared to traditional emergency services.

The ring does not replace Triple Zero (000). Participants are told to always call 000 for serious emergencies. The ring provides additional first-contact from community members during the minutes before professional help arrives.

The study will run for 12 months with approximately 500 adult participants.

Research Question

Does a wearable-initiated community proximity alert system reduce median time-to-first-contact in emergency events, compared to centralised emergency dispatch alone, when operating as a supplementary layer alongside existing Triple Zero services?

Study Design

Type: Prospective observational cohort with historical control comparison.

Design: Single-arm deployment of the intervention (NFC ring + mesh alert network) with comparison against historical EMS response time data for the same geographic area from the preceding 24 months.

Rationale for single-arm: Randomisation (giving some participants rings and not others in the same community) would compromise the mesh density that the system requires to function. The system's value depends on network density; a randomised design would test the system under artificially degraded conditions.

14.2 Ethical Framework

This study is governed by two overlapping ethical frameworks: the National Statement on Ethical Conduct in Human Research (NHMRC, 2023 update) and the OMXUS Principles — the non-negotiable ethical architecture of the system being tested.

National Statement Compliance

1. **Research Merit and Integrity:** The evidence synthesis (Part I of this thesis) establishes that the research question is meaningful, the methodology is appropriate, and the pilot is the minimum viable study to resolve empirical unknowns.
2. **Justice:** The benefits and burdens are fairly distributed. Participants are drawn from the community that would benefit from the system.
3. **Beneficence:** The expected benefit (faster emergency response, improved community cohesion) is proportionate to the identified risks.
4. **Respect for Persons:** Participation is voluntary, informed consent is comprehensive, withdrawal is unconditional.

OMXUS Principles Alignment

The OMXUS system is built on six hierarchically ordered principles. Three are absolute constraints; three are implementation requirements.

Principle 1: Cannot Affect Individual Freedom (Absolute). Responding to an alert is always voluntary. No social, reputational, or system-level penalty exists for non-response. This is not merely a design choice — it is an architectural constraint. The system records no per-person response/non-response data visible to others.

Principle 2: Non-Maleficence — Architectural, Not Promissory (Absolute). The system cannot be weaponised against its users. Not "we promise not to" — architecturally impossible. Alert broadcasts contain no personal identifier. Location is shared only during an active alert, only to devices within mesh range, and only for the alert duration. 15-minute session key rotation prevents temporal linkability. There is no continuous location tracking — the architecture does not support it.

Principle 3: Justice = Prevention Only (Absolute). The system's model of justice is prevention, not punishment. If misuse occurs, the response is education, support, or exclusion — not punishment, public shaming, or referral for sanction.

Principle 4: Transparent Accountability (Implementation). All study protocols, data collection instruments, and analysis plans are public.

Principle 5: Telemetry for Humans (Implementation). Your data works FOR you. This is not surveillance OF you — it is intelligence ABOUT you, owned BY you, serving YOU.

Principle 6: Zero Effort, Enjoyable, Instant Rewards (Implementation). The NFC ring is designed for zero-friction activation (gross motor triple-tap). Participation burden is minimised. The system provides immediate feedback.

Conflicts of Interest

The research team includes members of the OMXUS project, which developed the technology being tested. This is a substantive conflict managed through:

1. Independent DSMB with no financial relationship to OMXUS
2. Pre-registered analysis plan with ANZCTR
3. Open data: complete de-identified dataset published regardless of results
4. Independent statistical analysis by statistician with no involvement in technology development
5. Negative results commitment: results published regardless of outcome
6. No commercialisation during study
7. Full investigator disclosure in all publications

14.3 Study Design and Methods

Setting

One suburban community in [STATE], Australia. Site selection criteria: - Population density 1,000-5,000 per km^2 - Defined geographic boundaries - Existing community organisations willing to partner - Historical EMS response data available - Not currently participating in another community safety intervention

Participants

Target: n = 500 participants.

Justification: At 500 participants over approximately 4 km^2, the network achieves a raw density of 125/km^2. With an estimated willingness factor of w = 0.20, effective responder density is 25/km^2, yielding an expected nearest-responder distance of ~100m and estimated response time of 1-3 minutes. Based on Australian emergency incidence data (~1 per 4,600 people per day), the pilot expects 40-100 alert events over 12 months — sufficient for descriptive statistics on response times and acceptance rates.

Inclusion Criteria: - Adults aged 18 years or older - Resident or regularly present (>=4 days/week) in the study area - Owns or has regular access to a compatible smartphone (Android 10+ or iOS 15+) - Able to provide informed consent - Willing to wear NFC ring and have app installed for study duration

Exclusion Criteria: - Currently subject to an AVO or similar protective order as the respondent - Cognitive impairment that prevents informed consent - Participation would pose a safety risk (e.g., active psychosis, known violent behaviour)

Intervention

Each participant receives: - One NFC ring (passive, no battery, waterproof, multiple sizes) - Smartphone app (Android/iOS) - 30-minute onboarding session

Training covers: 1. How to activate an alert (triple tap on ring, or in-app button) 2. How to cancel a false activation 3. What to do when receiving an alert 4. Explicit instruction: always call 000 for serious emergencies 5. Explicit instruction: never enter a dangerous situation 6. Explicit instruction: you are never required to respond 7. Protocol for encountering a deceased person 8. Protocol for mental health crisis situations 9. How to withdraw from the study 10. Privacy: what data is collected and how 11. Mandatory reporting obligations

Outcomes

Primary Outcome: Median time from alert activation to first physical contact by a community responder (minutes), compared with historical median EMS response time.

Secondary Outcomes: 1. Alert acceptance rate 2. Acceptance rate by emergency type, time of day, responder demographics 3. False alarm rate 4. Change in acceptance rate over time (fatigue measure) 5. Mesh network alert propagation latency (seconds) 6. NFC ring vs. in-app activation rate 7. Responder adverse events 8. Participant-reported perceived safety changes 9. Participant-reported social cohesion changes

Statistical Analysis

Pre-registered with ANZCTR. Primary analysis: Mann-Whitney U test comparing community responder contact time with historical EMS response time. Secondary: logistic regression, time-series analysis, paired survey comparisons, thematic analysis of focus groups.

14.4 Risk Assessment

Risk	Likelihood	Severity	Mitigation
Responder injury	Low	High	Training prohibits entering dangerous situations. Voluntary response. Insurance.
False alarm fatigue	Moderate	Low	Rate limiting (3/day). Confirmation gesture. Automatic suppression.
System misuse (stalking)	Low	High	No continuous tracking. No identity in alerts. 15-min key rotation. AVO holders excluded.
Reliance on system over 000	Low	High	Training emphasises 000. App displays “Call 000” on every alert.

Risk	Likelihood	Severity	Mitigation
DV perpetrator locates victim	Low	High	Silent activation. No identity in broadcast. Coarse location only. AVO exclusion.
Responder psychological distress	Moderate	Moderate	Post-incident debrief within 48h. Referral. Opt-out mechanism.
Encounter with deceased person	Low	High	Training protocol. Critical incident support within 24h.
Mental health crisis escalation	Low-Mod	Moderate	Training: be present, do not restrain, call 000/crisis line.
Data breach	Low	Moderate	XChaCha20-Poly1305 encryption. Minimal data. Auto-delete after 90 days.
System failure during emergency	Moderate	Moderate	System is supplementary. 000 remains available.
Community conflict over non-response	Low	Low	No penalty for non-response (architectural). Anonymised acknowledgments.
Social pressure to respond	Low-Mod	Moderate	Anonymised acknowledgments. No leaderboards. No gamification.
Police interaction at scene	Moderate	Moderate	Training: identify as community member, cooperate, do not provide device without warrant.

Overall Risk Assessment: Low-moderate. The system is supplementary to existing services. Non-maleficence is enforced by cryptographic protocol, not by policy promise. Individual freedom is preserved by the absence of coercive mechanisms, not by rules against their use.

14.5 Privacy and Data Protection

Data Minimisation: No continuous location tracking. Alert broadcasts contain no personal identifier. Location shared only during active alert, only to nearby devices. Session keys rotate every 15 minutes. Relay nodes cannot read alert content.

Data Sovereignty: Participants own their data. They can request export or deletion at any time. No participant data will be sold, licensed, or shared with third parties for commercial purposes.

Law Enforcement: Study alert data will not be voluntarily provided to law enforcement. Exception: valid court order or warrant.

14.6 Vulnerable Populations

Domestic and Family Violence Survivors: Silent activation mode. No identity in broadcast. AVO respondents excluded. DV-specific protocol (see Appendix F).

Children: Not study participants (minimum age 18). May be beneficiaries. Mandatory reporting obligations apply.

Aboriginal and Torres Strait Islander Peoples: Community consultation. Culturally appropriate materials. AIATSIS Code of Ethics observed.

People with Disability: Ring in multiple form factors. WCAG 2.1 AA app. Alternative activation methods. Not assumed to be only alerters.

People with Mental Health Conditions: Active psychosis excluded. Managed conditions not excluded. Additional support offered.

14.7 Timeline and Budget

Phase	Duration	Activities
Ethics and approvals	Months 1-3	HREC submission, site agreements, ANZCTR registration
Community engagement	Months 3-4	Community information sessions, partner agreements
Site preparation	Months 4-6	Emergency services MOU, hardware procurement, app testing
Recruitment	Months 6-9	Participant recruitment and onboarding (rolling)
Active deployment	Months 9-21	12-month active monitoring period
Data analysis	Months 21-24	Statistical and qualitative analysis
Reporting	Months 24-26	Final report, publications, community debrief

Budget: AUD $170,950 (including 10% contingency). Major items: Research assistant $65,000, App development $25,000, Clinical psychologist $15,000, Independent statistician $10,000, NFC rings $7,500, Community engagement $8,000.

References

American Bar Association. (2021). How you start is how you finish? The slave patrol and Jim Crow origins of policing. *Human Rights Magazine*, 46(2).

American Heart Association. (2020). CPR and ECC guidelines: Executive summary. *Circulation*, 142(16 Suppl 2).

Australian Bureau of Statistics. (2022). Causes of Death, Australia. ABS.

Australian Bureau of Statistics. (2025). Corrective services, Australia, December quarter 2025. ABS.

Australian Institute of Health and Welfare. (2023). Health expenditure Australia 2022-23. AIHW.

Blackwell, T. H. et al. (2002). Response time effectiveness: Comparison of response time and survival in an urban emergency medical services system. *Academic Emergency Medicine*, 9(4), 288-295.

Bond, C. F. & DePaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review*, 10(3), 214-234.

Bor, J., Venkataramani, A. S., Williams, D. R., & Tsai, A. C. (2018). Police killings and their spillover effects on the mental health of Black Americans. *The Lancet*, 392(10144), 302-310.

Braga, A. A., Weisburd, D., & Turchan, B. (2019). Focused deterrence strategies effects on crime: A systematic review. *Campbell Systematic Reviews*, 15(3), e1051.

Braun, V. & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101.

Brooks, S. C., Simmons, G., Worthington, H., Bobrow, B. J., & Morrison, L. J. (2016). North American public opinion survey on the acceptability of crowdsourcing basic life support for out-of-hospital cardiac arrest with the PulsePoint mobile phone app. *JMIR mHealth and uHealth*, 4(2), e60.

Brooks, S. C. et al. (2022). Implementation of the PulsePoint smartphone application for crowdsourcing bystander resuscitation. *Journal of Clinical Medicine*, 15(1), 5.

Butts, J. A., Roman, C. G., Bostwick, L., & Porter, J. R. (2015). Cure Violence: A public health model to reduce gun violence. *Annual Review of Public Health*, 36, 39-53.

Centers for Disease Control and Prevention. (2024). The value of community paramedicine. CDC EMS and Community Paramedicine.

Chalfin, A., & McCrary, J. (2018). Are US cities underpoliced? Theory and evidence. *Review of Economics and Statistics*, 100(1), 167-186.

Chan, J. et al. (2019). Community paramedicine: A systematic review of program descriptions and training. *CJEM*, 21(6), 749-761.

Chartres, N. et al. (2016). Association of industry sponsorship with outcomes of nutrition studies. *JAMA Internal Medicine*.

Chassaing, B. et al. (2015). Dietary emulsifiers impact the mouse gut microbiota. *Nature*.

Choi, B. Y. et al. (2023). Community paramedicine intervention reduces hospital readmission and emergency department utilization. *Prehospital Emergency Care*, 27(7), 885-892.

Cohen, L. E. & Felson, M. (1979). Social change and crime rate trends. *American Sociological Review*, 44(4), 588-608.

Cure Violence Global. (2024). The effectiveness of the Cure Violence approach. CVG Impact Report.

Darley, J. M. & Latane, B. (1968). Bystander intervention in emergencies. *Journal of Personality and Social Psychology*, 8(4), 377-383.

Davis, A. Y. (2003). *Are Prisons Obsolete?* Seven Stories Press.

Disputes Centre Australia. (2024). Traditional Indigenous conflict resolution methods in Australia.

Eastridge, B. J. et al. (2012). Death on the battlefield (2001-2011): Implications for the future of combat casualty care. *Journal of Trauma and Acute Care Surgery*, 73(6), S431-S437.

Emsley, C. (2009). *The Great British Bobby: A History of British Policing from the 18th Century to the Present*. Quercus.

European Resuscitation Council. (2021). ERC Guidelines. *Resuscitation*, 161.

Evans, W. N. & Owens, E. G. (2007). COPS and crime. *Journal of Public Economics*, 91(1-2), 181-201.

Fischer, P. et al. (2011). The bystander-effect: A meta-analytic review. *Psychological Bulletin*, 137(4), 517-537.

Frazier, A. et al. (2021). Every one-minute delay in EMS on-scene resuscitation after out-of-hospital pediatric cardiac arrest lowers ROSC by 5%. *Resuscitation*, 166, 103-110.

Fulham, L., Blais, J., Rugge, T., & Schultheis, E. A. (2025). The effectiveness of restorative justice programs: A meta-analysis. *European Journal of Criminology*, 22(1).

GBD Police Violence Collaborators. (2021). Fatal police violence by race and state in the USA, 1980-2019. *The Lancet*, 398(10307), 1239-1255.

Gilmore, R. W. (2007). *Golden Gulag: Prisons, Surplus, Crisis, and Opposition in Globalizing California*. University of California Press.

Hadden, S. E. (2001). *Slave Patrols: Law and Violence in Virginia and the Carolinas*. Harvard University Press.

Harvard Political Review. (2021). Police unions are anti-labor.

Hatzolah Melbourne. (2024). Why call Hatzolah. https://hatzolah.org.au/why-call-hatzolah/

Heckman, J. J., Moon, S. H., Pinto, R., Savelyev, P. A., & Yavitz, A. (2010). The rate of return to the HighScope Perry Preschool Program. *Journal of Public Economics*, 94(1-2), 114-128.

Holmen, J. et al. (2021). Shortening ambulance response time increases survival in out-of-hospital cardiac arrest. *Journal of the American Heart Association*, 10(20), e020048.

IARC/WHO. (2015). IARC Monographs evaluate consumption of red meat and processed meat. Press release, October 26, 2015.

International Association of Fire and Rescue Services (CTIF). (2023). Global firefighting statistics report.

Kaba, M. (2021). *We Do This 'Til We Free Us: Abolitionist Organizing and Transforming Justice*. Haymarket Books.

Kadish, A. H. et al. (2007). Hatzolah emergency medical responder service: To save a life. *Medical Journal of Australia*, 186(12), 639-642.

Kearns, C. E. et al. (2016). Sugar industry and coronary heart disease research. *JAMA Internal Medicine*.

Larsen, M. P., Eisenberg, M. S., Cummins, R. O., & Hallstrom, A. P. (1993). Predicting survival from out-of-hospital cardiac arrest. *Annals of Emergency Medicine*, 22(11), 1652-1658.

Latimer, J., Dowden, C., & Muise, D. (2005). The effectiveness of restorative justice practices: A meta-analysis. *Prison Journal*, 85(2), 127-144.

Legewie, J. & Fagan, J. (2019). Aggressive policing and the educational performance of minority youth. *American Sociological Review*, 84(2), 220-247.

Lentz, S. A. & Chaires, R. H. (2007). The invention of Peel's principles. *Journal of Criminal Justice*, 35(1), 69-79.

Mapping Police Violence. (2025). 2024 police violence report. https://policeviolencereport.org/

Marmot, M. (2005). Social determinants of health inequalities. *The Lancet*, 365(9464), 1099-1104.

McCann, D. et al. (2007). Food additives and hyperactive behaviour in 3-year-old and 8/9-year-old children. *The Lancet*.

Medical Journal of Australia. (2025). Association of smartphone-activated volunteer responders with survival after out-of-hospital cardiac arrest. *MJA*, 222(10).

NAACP. (2024). The origins of modern day policing. https://naacp.org/

National Association for Community Mediation. (2023). Community mediation outcomes data.

National Law Enforcement Officers Memorial Fund. (2024). Slave patrols: An early form of American policing.

National Low Income Housing Coalition. (2024). The evidence is clear: Housing First works. NLIHC Brief.

Nettelbeck, A. & Smandych, R. (2010). Policing Indigenous peoples on two colonial frontiers. *Australian and New Zealand Journal of Criminology*, 43(2), 356-375.

NFPA 1710. (2023). Standard for the Organization and Deployment of Fire Suppression Operations.

NHMRC. (2023). National Statement on Ethical Conduct in Human Research (2023 update). Canberra.

Ornato, J. P. et al. (2012). Impact of dispatcher-assisted bystander CPR. *Resuscitation*, 83(7).

Patterson, D. G. et al. (2023). Evaluating the impact of a mobile integrated health-community paramedicine program. *BMC Health Services Research*.

Perkins, K. B. (1990). Volunteer fire departments: Community integration, autonomy, and survival. *Human Organization*, 49(4), 342-348.

Picard-Fritsche, S. & Cerniglia, L. (2013). Testing a public health approach to gun violence: Crown Heights. Center for Court Innovation.

PMC. (2022). Smartphone dispatch of lay volunteers to retrieve AEDs and respond to cardiac arrests. Randomized Clinical Trial.

Police Funding Database. (2024). https://policefundingdatabase.org/

Productivity Commission (Australia). (2026). Report on government services 2026: Police services.

PulsePoint Foundation. Community CPR and AED awareness program data. https://www.pulsepoint.org/

Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster.

RACGP. Preventative health investment returns. Royal Australian College of General Practitioners.

RAND Corporation. (1994). Controlling cocaine: Supply versus demand programs. RAND Drug Policy Research Center.

Reichel, P. L. (1988). Southern slave patrols as a transitional police type. *American Journal of Police*, 7(2), 51-77.

Ringh, M. et al. (2015). Mobile-phone dispatch of laypersons for CPR in out-of-hospital cardiac arrest. *New England Journal of Medicine*, 372(24), 2316-2325.

Roman, C. G. et al. (2018). Philadelphia CeaseFire: Findings from an outcome impact assessment. Temple University.

Sampson, R. J., Raudenbush, S. W. & Earls, F. (1997). Neighborhoods and violent crime. *Science*, 277(5328), 918-924.

Saver, J. L. (2006). Time is brain — quantified. *Stroke*, 37(1), 263-266.

Sentencing Project. (2024). One in five: Disparities in crime and policing.

Sewell, A. A. & Jefferson, K. A. (2016). Collateral damage: The health effects of invasive police encounters. *Journal of Urban Health*, 93(1), 42-67.

Shapiro, G. K. et al. (2021). Co-responding police-mental health programs: A review. *Administration and Policy in Mental Health*, 48, 1-18.

Sharkey, P., Torrats-Espinosa, G., & Takyar, D. (2017). Community and the crime decline. *American Sociological Review*, 82(6), 1214-1240.

Sherman, L. W. & Strang, H. (2007). *Restorative Justice: The Evidence*. Smith Institute.

Siegel, K. R. et al. (2016). Association of higher consumption of foods derived from subsidized commodities with adverse cardiometabolic risk. *JAMA Internal Medicine*.

Simons, F. E. (2011). Anaphylaxis: Recent advances in assessment and treatment. *Journal of Allergy and Clinical Immunology*, 127(3), 589-593.

Simpson, P. (2008). Hatzolah emergency medical responder service: To save a life. In *Volunteer Emergency Services in the 21st Century*.

Skogan, W. G., Hartnett, S. M., Bump, N., & Dubois, J. (2009). Evaluation of CeaseFire-Chicago. Northwestern University.

Smith, C. M. et al. (2020). Community First Responders in Real Emergencies. *Resuscitation*, 155, 152-159.

Stinson, P. M. (2020). Criminology explains police violence. *Contexts*, 19(2), 34-39.

Turner, K. B., Giacopassi, D., & Vandiver, M. (2006). Ignoring the past: Coverage of slavery and slave patrols in criminal justice texts. *Journal of Criminal Justice Education*, 17(1), 181-195.

Umbreit, M. S., Coates, R. B., & Vos, B. (2001). The impact of victim-offender mediation. *Federal Probation*, 65(3), 29-35.

United Hatzalah. (2024). Mission and impact. https://israelrescue.org/mission/

U.S. Bureau of Economic Analysis. (2024). Government current expenditures: Public order and safety: Police. FRED Economic Data.

Vitale, A. S. (2017). *The End of Policing.* Verso Books.

Webster, D. W. et al. (2012). Evaluation of Baltimore's Safe Streets program. Johns Hopkins Bloomberg School of Public Health.

Weisburd, D. et al. (2016). Place matters: Criminology for the twenty-first century. Cambridge University Press.

White Bird Clinic. (2024). What is CAHOOTS? https://whitebirdclinic.org/what-is-cahoots/

World Health Organization. Between 30-50% of all cancer cases are preventable. WHO Fact Sheet.

Yan, S. et al. (2020). The global survival rate among adult out-of-hospital cardiac arrest patients who received cardiopulmonary resuscitation. *Critical Care*, 24(1), 61.

Appendix A: Participant Information Sheet and Consent Form

A.1 Participant Information Sheet

Study Title: Civic Proximity Response Pilot Study **Ethics Approval Number:** [TO BE ASSIGNED] **Principal Investigator:** [NAME], [INSTITUTION] **Contact:** [PHONE] / [EMAIL]

What is this study about? We are testing a new way to help people in emergencies get help faster. You will receive a special ring and a phone app. If you or someone near you has an emergency, you can tap the ring to alert nearby participants. If someone near you taps their ring, you will get a notification and can choose to go help.

This does not replace calling 000. You should always call 000 for serious emergencies. The ring is an additional way to get immediate help from people who are already nearby.

What will I need to do? - Attend a 30-minute onboarding session - Wear the NFC ring (or carry it) - Have the study app installed on your phone - Complete 3 short surveys (start, 6 months, 12 months) - Optionally: respond to emergency alerts from nearby participants

Do I have to respond to alerts? No. You never have to respond. Receiving an alert does not create any obligation. You can ignore any alert for any reason. There is no penalty, record, or consequence for not responding. The system does not track whether you respond or not.

What are the risks? - If you choose to respond, you may encounter an upsetting situation. You are trained to stay safe and never enter a dangerous situation. - You may experience psychological distress after responding to a traumatic event. Free, confidential support is available. - You may

receive false alarms. The system limits these. - Your approximate location is shared with nearby participants only during an active alert you trigger. There is no tracking at other times.

What are the benefits? - You may receive faster help in an emergency from people near you - You contribute to research that could improve emergency response for everyone - You keep the NFC ring after the study

How is my privacy protected? - No continuous location tracking - Location only shared during an active alert you trigger, only to nearby devices - All data encrypted - Alert logs automatically deleted after 90 days - Research data is de-identified - You can request deletion at any time - Your data will not be given to police, employers, insurers, or anyone else (except where required by law)

Can I withdraw? Yes. At any time. Your data deleted within 14 days. No penalty. No reason required.

Support services: - 1800RESPECT (1800 737 732) — family and domestic violence - Lifeline (13 11 14) — crisis support - Beyond Blue (1300 22 46 36) — mental health

A.2 Consent Form

Study Title: Civic Proximity Response Pilot Study **Ethics Approval Number:** [TO BE ASSIGNED]

I have read and understood the Participant Information Sheet (Version 3.0, February 2026).

- ☐ I understand the purpose and procedures of this study.
- ☐ I understand that my participation is voluntary and I can withdraw at any time without penalty.
- ☐ I understand that responding to emergency alerts is voluntary and I am never required to respond.
- ☐ I understand that this system does not replace Triple Zero (000).
- ☐ I understand the risks described in the information sheet, including the possibility of encountering distressing situations.
- ☐ I understand that I should never enter a dangerous situation when responding to an alert.
- ☐ I understand how my data will be collected, stored, and protected.
- ☐ I understand that my data will not be provided to police except under court order.
- ☐ I understand the mandatory reporting obligations.
- ☐ I consent to the research team accessing de-identified alert event data.
- ☐ I have had the opportunity to ask questions.
- ☐ I consent to participate in this study.

Participant Name: _____ Date: _____

Participant Signature: _____

Researcher Name: _____ Date: _____

Researcher Signature: _____

Appendix B: Survey Instruments

B.1 Perceived Safety Scale (Adapted)

Participants rate agreement (1 = Strongly Disagree to 5 = Strongly Agree):

1. I feel safe walking in my neighbourhood during the day.
2. I feel safe walking in my neighbourhood at night.
3. If I had an emergency, someone nearby would help me.
4. I would feel comfortable asking a neighbour for help.
5. I trust the people in my neighbourhood.
6. I would intervene if I saw someone who needed help.
7. I believe my neighbourhood is becoming safer / staying the same / becoming less safe.

B.2 Collective Efficacy Scale (Adapted from Sampson et al., 1997)

Social cohesion subscale: 1. People around here are willing to help their neighbours. 2. This is a close-knit neighbourhood. 3. People in this neighbourhood can be trusted. 4. People in this neighbourhood generally do not get along with each other. (R) 5. People in this neighbourhood share the same values.

Informal social control subscale ("How likely is it that your neighbours would intervene if…"): 1. Children were showing disrespect to an adult. 2. A fight broke out in front of their house. 3. Someone was being assaulted. 4. Someone appeared to be having a medical emergency. 5. A neighbour's property was being vandalised.

B.3 Post-Incident Debrief Guide

Conducted within 48 hours. Semi-structured interview: 1. What happened? 2. How did you activate/receive the alert? 3. How long did it take for someone to arrive? 4. Did you also call 000? Which arrived first? 5. How did you feel during the event? 6. How do you feel now? 7. Wellbeing check: intrusive thoughts, difficulty sleeping, heightened anxiety, avoidance? 8. Is there anything about the system you would change? 9. Do you need any support? 10. If deceased person encountered: additional questions by clinical psychologist.

Appendix C: Technology Safety Specification

C.1 Cryptographic Primitives

Function	Algorithm	Purpose
Identity key (IK)	Ed25519	Long-lived device identity
Derived key (DK)	HKDF-SHA256	Epoch-specific key derivation
Session key (SK)	Rotated every 15 min	Prevents temporal linkability
Alert encryption	XChaCha20-Poly1305	Authenticated encryption

Message encoding	CBOR (RFC 8949)	Compact binary serialisation

C.2 Privacy Guarantees

1. **No identity in broadcasts.** SOS alert messages contain no personal identifier, name, phone number, or account reference.
2. **Session key rotation.** Every 15 minutes. Cannot link two alerts from the same device across epoch boundaries.
3. **Relay blindness.** Relay nodes process encrypted payloads they cannot decrypt.
4. **Rate limiting.** 3 SOS alerts per 24 hours per identity. 6 audible notifications per hour per helper device. Flood protection: >20 SOS in 5 minutes triggers degraded mode.
5. **Escalation thresholds.** Alerts escalate to wider relay network only if <2 acknowledgments within 45 seconds.

C.3 Threat Model

Adversary	Capability	Protection
Passive observer	Monitors BLE broadcasts	No identity; session key rotation
Malicious relay	Forwards/drops/modifies	E2E encryption; relay cannot read
Stalker	Tracks specific person	15-min key rotation; no linkability
DV perpetrator	Locates victim	No identity; coarse location; silent mode
Law enforcement	Requests alert data	Minimal data; auto-deletion; no central store
Research team	Access to infrastructure	Only de-identified aggregates

C.4 What the Technology Cannot Prevent

- A participant verbally disclosing that they triggered or responded to an alert
- Observation of another's phone receiving an alert in the same room
- Traffic analysis by a sophisticated adversary controlling many devices (mitigated but not eliminated)
- The ring itself being a visible signal of study participation

---

Appendix D: Emergency Services Integration Protocol

D.1 Memorandum of Understanding

Prior to study commencement, an MOU will be sought with the relevant ambulance service and local police. The MOU establishes:

1. Emergency services informed of the community response network
2. System is supplementary to 000

3. Participants trained to yield to professional responders
4. Historical EMS response time data requested for comparison
5. Incident coordination for debriefing purposes

D.2 Scene Protocol

When professional services arrive where a community responder is present: 1. Community responder identifies themselves as a neighbour/community member 2. Provides brief handover of observations and actions 3. Steps back and follows professional instructions 4. Does not leave until released

D.3 Police Investigation at Scene

- Responder cooperates as any member of the public would
- Does not disclose details about the alert system or other participants
- Does not hand over phone or device without a warrant
- Research team available 24/7 for police enquiries

Appendix E: Responder Safety and Wellbeing Protocol

E.1 Before Response: Training

1. **Your safety comes first.** If the situation appears dangerous, do not approach.
2. **You are not a first responder.** You have no duty to provide medical care or restrain anyone.
3. **"Observe and report" is a complete response.** Arriving, assessing, calling 000, and waiting is a full, valid, valued response.

E.2 During Response: Safety Rules

1. Do not enter a building if you cannot see what is inside
2. Do not approach a violent or aggressive person
3. Do not attempt to restrain anyone
4. Do not move an injured person unless in immediate danger
5. Do not provide interventions beyond basic first aid
6. If you feel unsafe at any point, leave immediately

E.3 After Response: Wellbeing Support

Routine Debrief: Within 48 hours, by research team member.

Critical Incident Response (death, serious injury, violence, child in danger, or responder distress): 1. Immediate: contacted within 4 hours 2. 24 hours: debrief with clinical psychologist 3. 7 days: follow-up welfare check 4. 30 days: second follow-up 5. Ongoing: referral to external services if needed

Opt-Out: At any time, disable responder function while retaining ability to trigger alerts. Temporarily or permanently. Without explanation.

Appendix F: Domestic and Family Violence Safety Protocol

F.1 Design Safeguards

1. **Silent activation.** No sound, no vibration, no screen change on alerter's device.
2. **No identity disclosure.** Alerts contain no personal identifier.
3. **Coarse location only.** Grid-cell level (~100m^2), not GPS coordinates.
4. **AVO exclusion.** Respondents to active protective orders excluded.
5. **Selective blocking.** Participants can request specific individuals not receive their alerts.

F.2 Onboarding

All participants receive DFV service information (1800RESPECT, state DV lines), explanation of silent mode, and private opportunity to discuss safety concerns with DFV-informed researcher.

F.3 Scenario Protocols

Responder arrives at DFV situation: Do not enter. Call 000. Remain at safe distance. Do not confront. Contact research team.

Participant discloses DFV: Provide service contacts. Do not intervene directly. Mandatory reporting if children at risk. Offer participation adjustment.

Perpetrator attempts to use system to locate victim: Architecture prevents this. If suspected misuse reported, immediate investigation. Perpetrator terminated from study.

Appendix G: Data Flow and Privacy Architecture

G.1 Data Categories

Data Type	Collected By	Stored	Retention	Access
Alert event logs	App (auto)	Device only (encrypted)	90 days, auto-deleted	Participant only
De-identified telemetry	App to research server	Institutional server (encrypted)	Study + 5 years	Research team, DSMB
Survey responses	Paper/online	Institutional server	Study + 5 years	Research team
Participant identifiers	Consent form	Locked filing cabinet / encrypted file	Study + 5 years	PI and RA only
Mesh network metadata	App (auto)	Aggregated on server	Study duration	Research team

Data Type	Collected By	Stored	Retention	Access
Post-incident debriefs	Research team	Institutional server (encrypted)	Study + 5 years	Research team

G.2 Separation of Concerns

1. **Alert data never reaches OMXUS.** Device-to-device only.
2. **Research data held by institution,** not OMXUS.
3. **Participant identifiers separated** from research data.
4. **Published data is aggregated.** No individual histories.

G.3 Participant Data Rights

1. Right to access
2. Right to correction
3. Right to deletion (within 14 days)
4. Right to export (machine-readable format)
5. Right to know (breach notification within 72 hours)

Appendix H: Community Engagement Plan

H.1 Pre-Study (Months 3-4)

1. At least 3 public community information sessions
2. Local council briefing
3. Emergency services briefing
4. Community advisory group (5-8 members, mix of participants and non-participants)
5. Aboriginal and Torres Strait Islander consultation where appropriate
6. CALD community outreach and translation services

H.2 During Study (Months 6-21)

1. Community advisory group meets quarterly
2. Six-month community update session
3. Mechanism for non-participants to raise concerns
4. Local media updates at PI's discretion

H.3 Post-Study (Months 24-26)

1. Community debrief session with findings
2. Plain language summary to all participants and partners
3. Community advisory group consulted on dissemination
4. If system continued, community vote on continuation

Appendix I: Coverage Density Calculations

Willingness (w)	90% Coverage (per km^2)	95% Coverage (per km^2)	Adoption needed (urban 20k/km^2)	Adoption needed (suburban 3k/km^2)
0.10	183	238	1.2%	7.9%
0.15	122	159	0.8%	5.3%
0.20	92	119	0.6%	4.0%
0.30	61	79	0.4%	2.6%
0.50	37	48	0.2%	1.6%

Appendix J: Cross-References — OMXUS Research Series

This thesis is Paper No. 13 in the OMXUS Research Series. It provides evidence for Conclusions #3 (The justice system punishes the wrong people), #4 (Community safety outperforms state policing), and #15 (Proximity-based knowledge is systematically ignored).

Direct Connections

Paper	Title	Connection to This Thesis
Paper 3	Prevention Over Punishment	Establishes the fiscal case for prevention over punishment. This thesis provides the operational model — Hatzolah's 2-4 minute response times versus 8-14 minutes for traditional EMS — proving that community-based infrastructure delivers superior outcomes at lower cost. Located: `content/research/prevention_over_punishment/`
Paper 4	Universal Basic Income	Addresses the economic desperation that drives much of what the justice system processes as crime. When survival needs are met, the demand for punitive policing drops, making community-first models not just possible but natural. Located: `content/research/labor_economics_22hr_week/`
Paper 5	Two Monkey Theory	Establishes that fairness norms are biologically encoded. Community-first safety models succeed because they operate within the social scale at which those fairness instincts function — where people know each other's names, not badge numbers. Located: `content/research/two_monkey_theory/`
Paper 11	Signal Inversion (Constructed Guilt)	Demonstrates that police interrogation methods systematically produce false confessions. This thesis traces that dysfunction to its historical root, showing that modern policing was designed for economic enforcement, not community safety, which explains why its credibility-assessment methods are inverted. Located: `content/research/constructed_guilt_signal_inversion/`

Supporting Research in the Series

Research Area	Location	Relevance
Bystander Effect	`content/research/bystander_effect/`	Psychological research on visibility, critical mass, and sympathy gradients that inform the response system design. The ring system is specifically designed to overcome diffusion of responsibility through personal, named alerts.
Direct Personal Alerts	`content/research/direct_personal_alerts/`	Research on direct alert mechanisms (PulsePoint, GoodSAM) that overcome diffusion of responsibility in emergencies. Provides the behavioural evidence for why "You, Sarah, are 47 metres away" works when "Someone call an ambulance" does not.
BLE Mesh Networking	`content/research/ble_mesh_networking/`	The technical mesh infrastructure that enables sub-second alert propagation. No cellular dependency, no internet requirement, no central point of failure or censorship.
Sybil Resistance / Physical Presence	`content/research/sybil_resistance_physical_presence/`	Physical-presence verification ensuring only real humans are in the responder network. Prevents gaming, bot attacks, and false identity in the mesh.
Consensus Distillation / Trust	`content/research/consensus_distillation_trust/`	Identity verification and trust networks that underpin the responder community. How you know the person who shows up is who they say they are.

Research Area	Location	Relevance
Drug Policy Reform	content/research/drug_policy_reform/	Portugal model evidence. 80% fewer overdose deaths through decriminalisation. Relevant because drug-related calls constitute a significant proportion of emergency service demand. Community-first response (CAHOOTS model) handles these calls without police.
Community Policing Alternatives	content/research/community_policing_alternatives/	Evaluation of alternative safety models including violence interruption, restorative justice, and community mediation.
Housing First	content/research/housing_first/	Evidence that stable housing reduces police contacts by 40%, arrests by 40%, and homelessness by 88%. The most cost-effective "policing" intervention is not policing at all — it is a house.
Play Deprivation	content/research/play_deprivation/	Evidence that play-deprived children develop antisocial behaviours. Relevant to Goal 12 (education redesign) and the prevention dividend: environments designed for human bodies (Goal 11: monkey bars at every bus stop) produce healthier, safer communities.
Food Toxicology / Safety	content/research/food_toxicology_safety/	Evidence base for Goal 10 (food contains only things proven safe) and Chapter 11 of this thesis (children at risk from dietary policy failures).

Research Area	Location	Relevance
Social Group Scaling	content/research/social_group_scaling/	Dunbar's 150 scaling is discredited (Lindenfors et al. 2021: CI of 2-520). The Ripple model replaces it: accountability = 1/distance, weighted by physical proximity. Community-first safety works because it operates on the proximity gradient — whoever is nearest responds.
Education / Prussian Model	content/research/education_prussian_model/	The Prussian model of education system produces populations that defer to authority rather than organise their own safety. Relevant to why community-first models face political resistance despite superior evidence.
Economic Servitude	content/research/economic_servitude/	The 40-hour work week eliminates the time required for community governance and mutual aid. Goal 2 (22-hour work week) is a prerequisite for Goal 13 (community emergency response) because you cannot be a community responder if you are at work 50 hours a week.

The Convergence

Every paper in this series proves every other. If policing was designed to protect property rather than people (this thesis), and if its credibility detection is systematically inverted (Paper 11), then the $237.7 billion spent annually on US policing is not a safety investment but a control investment — and the community-based alternatives documented here are what actual safety looks like.

If the work week consumes all available time for community participation (economic servitude research), and if communities deprived of social infrastructure produce worse health outcomes (social group scaling), then the 22-hour work week (Goal 2) is not a labour policy — it is a safety policy. It is the time that makes community emergency response possible.

If the education system produces compliance rather than initiative (Prussian model research), and if play deprivation produces antisocial behaviour (play deprivation research), then redesigning education around play and mastery (Goal 12) is not an education policy — it is a crime prevention

policy.

The fourteen goals are not fourteen separate proposals. They are one system. This thesis — Goal 13, the ring — is the acute intervention. The other thirteen goals are the chronic cure.

This thesis is part of the OMXUS Research Series on decentralised infrastructure for human safety and flourishing.

Correspondence: research@omxus.com