

SCENARIO 1 – DIVERSION OF A JOURNALIST'S AIR TAXI THROUGH DISCRETE STATE PRESSURE

Human Rights at Risk

- Freedom of expression
- Freedom of movement
- Right to peaceful assembly

Assumptions:

The year is 2035 and Advanced Air Mobility (AAM) services have become a routine mode of transportation in many large metropolitan regions. Most passenger-carrying eVTOL services operate without an onboard pilot. Instead, aircraft operate at high levels of automation under the supervision of certified remote pilots and operational controllers located in centralized operations centers. Remote pilots simultaneously supervise small numbers of aircraft and retain the authority to intervene in flight operations, modify routes, or execute diversions when required by operational, safety, or security considerations.

Governments possess broad regulatory oversight powers over aviation operators, including authority relating to public safety, airspace access, and national security. While most democratic states maintain strong legal safeguards for civil liberties, some governments have increasingly relied on informal pressure, administrative coercion, and public-private partnerships to influence private-sector actors without issuing formal legal orders. Air taxi operators depend heavily on government licences, route permissions, vertiport access agreements, and regulatory approvals, creating incentives to comply with government requests even when those requests fall outside established legal procedures.

In this scenario, the air taxi operator is a large, reputable company operating legally and in accordance with all aviation safety regulations. The flight itself is technically safe, the aircraft systems function correctly, and no cybersecurity breach occurs. The human rights violation emerges from decisions made by human actors operating within a complex political and commercial environment, where informal government pressure influences operational choices without explicit legal compulsion.

Scenario:

In 2035, air taxis have become an integral part of daily transportation across several major cities. Dozens of short-range eVTOL flights occur every day, connecting suburbs, business districts, airports, and regional urban centers. Because ground transportation congestion has worsened significantly over the previous decade, politicians, business leaders, journalists, and academics increasingly rely on AAM services to travel rapidly between events.

Elena is one of the country's most prominent investigative journalists. Over the previous three years, she has published a series of articles alleging corruption involving government procurement contracts, urban infrastructure projects, and the licensing of emerging transportation

technologies, including AAM aircraft. Although she has never been charged with any crime, government officials frequently describe her reporting as "irresponsible", "destabilizing", and contrary to the national interest.

In recent months, Elena has become an increasingly visible public figure. She has been invited to deliver a keynote address at a conference on democratic governance and technological accountability hosted by a consortium of universities and civil society organizations. The conference will be livestreamed internationally and is expected to receive substantial media coverage.

On the morning of the event, Elena booked a scheduled air taxi flight operated by SkyLift Corp., one of the country's largest AAM service providers. The journey from her home city to the conference venue would take approximately forty minutes by air, compared with over three hours by road due to planned highway closures and severe congestion.

The aircraft departed normally. As is standard practice, no pilot is physically onboard. Instead, the aircraft operates autonomously under the supervision of a certified remote pilot located at SkyLift Corp.'s regional operations center, supported by an automated network management system and a team of operational supervisors. Weather conditions are favorable, and traffic density is moderate.

Immediately after take-off, SkyLift Corp.'s government liaison office receives an urgent communication from a senior official within the Ministry of Transport. No written order is issued. Instead, the official expresses concern that the passenger on board may be travelling to participate in activities that could "threaten public order and national stability". The official reminds SkyLift Corp. personnel that the company currently has several pending applications for expanded route authorisations and additional vertiport operating licences awaiting government approval. The message is not framed as a direct instruction. Nevertheless, senior company managers understand the implications immediately.

An emergency conference call is initiated between the airline operations manager, legal counsel, the head of remote operations, and a government relations executive. Several employees express concern that preventing the passenger from reaching her destination could violate company policy and international law. Others point out that refusal to cooperate could jeopardize the company's commercial future, potentially affecting thousands of employees and millions of dollars of investment. After a brief discussion, the operations manager authorises a diversion.

The remote pilot supervising Elena's aircraft receives instructions from senior management that an operational security concern requires the aircraft to divert to an alternate vertiport. The remote pilot requests clarification regarding the nature of the threat but is informed that the matter concerns a sensitive government security issue and that immediate compliance is expected. Although uncomfortable with the instruction, the pilot recognises that company procedures require compliance with operational directives issued by the command chain.

The aircraft's flight management system is reprogrammed remotely. Elena receives an automated notification informing her that the aircraft is being diverted because of "temporary operational constraints." When she attempts to contact customer support through the passenger communications interface, she receives only generic responses indicating that further information is unavailable.

The aircraft lands at a small suburban vertiport approximately forty kilometers from the conference venue. Upon landing, Elena discovers that no onward air services are available. Ground transportation options are limited due to a regional transport strike. By the time she reaches the conference venue several hours later, the event has concluded. International media outlets have begun reporting that she failed to appear.

In the following days, government officials publicly denied any involvement in the incident. SkyLift Corp. issues a statement citing "operational considerations" and declines to provide further details. Internally, several employees involved in the diversion express concern about what occurred. The remote pilot who executed the diversion files an anonymous ethics complaint but subsequently withdraws it after colleagues warn that challenging management decisions could damage future career prospects.

Over the following months, investigative reporting by a human rights NGO reveals a pattern of similar incidents involving political activists, opposition politicians, labor organizers, and human rights advocates. None of the affected individuals had been formally arrested, charged, or legally prohibited from travelling. Instead, authorities had repeatedly relied on informal communications with transport operators, exploiting commercial dependencies and regulatory leverage to restrict movement indirectly.

The resulting public inquiry concludes that no aviation safety regulations were violated and that all aircraft operated within their certified operational envelopes. However, the inquiry also determines that existing legal frameworks had failed to account for the ways in which highly automated, remotely operated mobility systems could enable subtle forms of political repression without requiring explicit state coercion, physical intervention, or technical compromise of the aircraft itself.

Preventative Actions

This scenario illustrates that human rights risks in Advanced Air Mobility may emerge not only from technological failures, but also from institutional relationships, governance structures, and economic incentives. The central challenge is preventing informal political influence from being exercised through private transportation providers.

One possible approach is to establish explicit legal protections prohibiting transport operators from restricting passenger movement except under clearly defined legal authorities. Any government request to delay, divert, or deny transport services for public order or national security reasons could be required to follow formal procedures, including written authorization, judicial

oversight, and independent review mechanisms. Such requirements would reduce opportunities for informal coercion and create an auditable record of government intervention.

A second approach would involve strengthening corporate human rights governance. AAM operators could be required to conduct human rights due diligence assessments. These assessments would identify situations where commercial incentives, licensing dependencies, or government relationships might create risks of human rights abuse. Operators could establish internal escalation procedures requiring review by independent ethics or human rights officers before passenger diversions or service denials are authorized.

A third approach would focus on transparency and accountability mechanisms. Operators might be required to maintain secure records of all government requests affecting passenger operations, including informal communications. Independent regulators or ombuds institutions could periodically review these records to identify patterns of political interference. Transparency reporting requirements, analogous to those used in the telecommunications sector, could further deter misuse.

Technical and operational safeguards could also reduce risks. Flight management systems could require explicit classification of the basis for operational diversions, creating a documented distinction between safety-related, operational, and governmental interventions. Decision-support systems might generate alerts when diversion decisions involve politically exposed persons, journalists, human rights defenders, or other individuals at heightened risk of abuse, triggering additional review procedures.

Finally, the broader AAM ecosystem, including regulators, operators, vertiport providers, insurers, and manufacturers, could incorporate human rights considerations into safety management systems and operational culture. Just as aviation developed a robust safety culture through incident reporting, learning systems, and root cause investigation, the AAM sector could develop analogous mechanisms for identifying, reporting, and mitigating emerging human rights risks before they become systemic.

Observations

The following should be considered:

- Human rights impact assessments for AAM operators and infrastructure providers.
- Mandatory procedures governing government requests affecting passenger transport operations.
- Independent oversight and auditing of operational diversions and passenger service denials.
- Confidential, anonymous reporting mechanisms for employees experiencing political or commercial pressure and the risk of reprisal.
- Human rights training for pilots, operations personnel, and corporate leadership.

ATPTI Case Studies

- Transparency reporting regarding government requests for passenger movement restrictions.
- Development of industry standards addressing political interference in AAM operations.
- Integration of human rights risk management into existing aviation safety management systems.
- International standards clarifying the responsibilities of AAM operators under the UN Guiding Principles on Business and Human Rights.

SCENARIO 2 – FAILURE OF AUTONOMOUS SYSTEMS TO DETECT HUMAN TRAFFICKING

Human Rights at Risk

- Freedom from slavery, servitude, and forced labour
- Freedom from human trafficking and exploitation
- Liberty and security of the person

Assumptions:

The year is 2047 and fully autonomous passenger-carrying eCTOL services have become operational in several jurisdictions. To compensate for the absence of onboard crew, aircraft companies try to minimize their liability by offering security systems on aircraft. These systems combine biometric identity verification, behavioral analytics, computer vision, and artificial intelligence-based anomaly detection to identify potential safety, security, and welfare concerns.

The airplane operator in this scenario has developed an advanced safety platform using cameras at airports, within passenger cabins, and throughout passenger processing areas. The system has been trained using large datasets of human behavior and is intended to identify indicators associated with drug intoxication and medical emergencies, two of the major safety risks in airplanes. The system is supported by a small number of remote officers who monitor automated alerts but do not routinely review all passenger interactions.

The operator complies with all applicable regulations and industry standards. The aircraft itself functions correctly, and all technical systems operate as designed. The human rights violation results not from technical malfunction, but from the goals the safety monitoring system was created to achieve.

Scenario

By 2047, autonomous eCTOL airplanes have become a common form of transportation within several large metropolitan regions. The elimination of onboard pilots and crew has significantly reduced operating costs, allowing operators to offer faster, low-cost transportation services connecting airports.

One of the companies with the largest market share in the sector is SkyLift Corp., which markets itself as the safest autonomous passenger transport provider in the world. A key component of its business model is its proprietary "Guardian" safeguarding platform, an AI-enabled passenger protection system designed to verify the identity of passengers and detect medical emergencies and drug intoxication.

The Guardian system performs multiple functions. Before boarding, passengers undergo biometric identity verification to ensure that only customers who paid for a ticket are boarding the flight. Cameras located throughout vertiports continuously analyze passenger behavior, facial

expressions, vocal characteristics, and body language. During flight, cabin cameras and environmental sensors continue to monitor passengers for signs of medical distress, violence, and drug intoxication.

Promotional material produced by the company claims that the system can identify indicators of unsafe situations with greater accuracy than human operators. These capabilities have become a major selling point for regulators and the public, particularly because autonomous aircraft no longer contain onboard personnel capable of intervening directly.

On a busy Friday afternoon in 2047, a man named Daniel books an autonomous aircraft across borders. Travelling with him is a sixteen-year-old woman named Sofia. Unknown to SkyLift Corp. personnel, Sofia is a victim of human trafficking who has been subjected to prolonged psychological coercion, threats of violence, and financial exploitation.

As they arrive at the departure airport, Sofia displays several behaviors that would traditionally concern experienced transport personnel had they been present. She speaks very little and appears anxious and visibly distressed. However, the Guardian system has not been trained to detect human trafficking, so it did not pick up on these behaviors.

Meanwhile, Daniel behaves calmly and cooperatively throughout the boarding process. His biometric credentials verify successfully, and he has no criminal record or previous safeguarding alerts associated with his account.

During the pre-boarding process, Sofia briefly raises her hand and performs a gesture that had, several years earlier, become internationally recognized as a covert signal indicating that an individual requires help while under coercion. However, the AI was not trained to recognize the gesture.

During the ninety-minute flight, the onboard safeguarding system continues to monitor passenger behavior. Sofia remains visibly uncomfortable and avoids interaction with Daniel. However, to minimize data storage, the system did not save any recording from the on-flight cameras because it did not pick up on any concerning behaviors.

Upon arrival, the passengers leave the airport without incident. Several months later, following an unrelated criminal investigation, law enforcement authorities uncovered an organized trafficking network operating across multiple regions using autonomous transportation services. Investigators discovered that traffickers had deliberately selected SkyLift Corp. because its automated safeguarding systems were not trained to identify or report coercive behavior.

A subsequent investigation by a human rights organization reveals that multiple trafficking victims had previously travelled on SkyLift Corp. services. In several cases, victims showed signs of human trafficking or even tried to alert the automated safety system, but their signals were not detected. Because human safeguarding personnel only reviewed cases that the automated safety system flagged, no intervention had occurred.

Public reaction is immediate and severe. Regulators, human rights organizations, and victims' advocates criticize the company's reliance on automated behavioral classification and question whether commercial pressures to reduce staffing levels had created systemic safeguarding vulnerabilities.

The resulting inquiry concludes that the safeguarding system had functioned exactly as designed. However, the inquiry also determines that the design assumptions underlying the system failed to account for instances of human trafficking with no obvious violence on board. The investigation further concludes that the absence of meaningful human oversight had transformed a tool intended to protect vulnerable individuals into a source of false reassurance.

Preventative Actions

This scenario illustrates the limitations of relying exclusively on artificial intelligence to identify complex human rights abuses in real time. Human trafficking, coercive control, and exploitation frequently involve subtle behavioral indicators that are difficult to distinguish from ordinary human behavior.

One possible approach is the implementation of hybrid human-AI safeguarding systems. Rather than treating artificial intelligence as a replacement for human judgement, operators could employ AI systems as decision-support tools that identify patterns requiring review by trained human safeguarding specialists. Artificial intelligence may excel at recognizing certain observable indicators, including repeated travel patterns, unusual booking behavior, known hand signals requesting assistance, or combinations of behavioral cues occurring over time. Human operators, however, may be better equipped to recognize contextual information, interpersonal dynamics, emotional distress, and subtle indicators of coercion. This inherently human skill is known as intuition.

A second approach would involve developing multiple independent safeguarding channels. Passengers experiencing coercion or exploitation could be provided with several methods for requesting assistance without alerting accompanying individuals. These might include recognized hand signals, silent digital distress functions integrated into passenger interfaces, coded verbal requests, biometric stress indicators, or physical emergency activation mechanisms within the cabin environment. Multiple channels reduce reliance on any single safeguarding method and provide redundancy when individual mechanisms fail.

A third approach would involve recognizing the limitations of behavioral prediction systems altogether. Rather than requiring AI systems to determine whether trafficking is occurring, safeguarding systems could instead identify situations involving elevated uncertainty or vulnerability and escalate those cases for human assessment. This approach accepts that many forms of exploitation cannot be reliably detected through behavioral analysis alone.

Organizational safeguards may also be necessary. Operators could be required to employ trained safeguarding specialists with expertise in human trafficking, coercive control, trauma

responses, and vulnerable populations. These specialists could review AI-generated alerts, conduct periodic audits of missed cases, and continuously evaluate whether system performance varies across demographic, cultural, or social groups.

Finally, regulators and operators should avoid creating unrealistic expectations regarding the capabilities of artificial intelligence. Claims that automated systems can reliably identify human trafficking or coercion may create dangerous overconfidence among operators, regulators, and the public. Safeguarding technologies should be understood as imperfect tools that support, rather than replace, human responsibility and professional judgement.

Observations

The following should be considered:

- Development of hybrid human-AI safeguarding systems combining automated detection with trained human review.
- Research into the reliability and limitations of behavioral indicators associated with human trafficking and coercive control.
- Recognition and standardization of covert distress signals, including hand gestures and digital assistance requests.
- Multiple independent channels for passengers to request assistance without alerting traffickers or coercive companions.
- Scenario-based training for remote safeguarding personnel involving trafficking, coercion, and exploitation scenarios.
- Periodic auditing of false negative and false positive safeguarding outcomes.
- Independent human rights impact assessments of behavioral surveillance and safeguarding technologies.
- Evaluation of the effects of commercial incentives on safeguarding staffing levels and intervention thresholds.
- Development of international standards governing the use of artificial intelligence in passenger safeguarding and anti-trafficking operations.
- Clear allocation of responsibility and accountability between operators, software developers, safeguarding personnel, and regulators when automated safeguarding systems fail.

SCENARIO 3 – AUTONOMOUS AERIAL SURVEILLANCE AND REMOTE LAW ENFORCEMENT OPERATIONS

Human Rights at Risk

Right to privacy

Freedom from arbitrary arrest and detention

Assumptions:

The year is 2029 and autonomous aerial systems have become widely integrated into urban public safety operations. Police agencies routinely employ networks of autonomous drones for traffic management, emergency response, crime prevention, and public order monitoring. These systems operate continuously using artificial intelligence, computer vision, biometric identification, predictive analytics, and automated decision-support systems.

In response to concerns about police officer safety, several jurisdictions have authorized the use of remotely operated and autonomous law enforcement drones to conduct certain policing functions previously performed by officers in person. These functions include surveillance, suspect tracking, public order management, and, under limited circumstances, the execution of arrests. The official justification for these policies is that they reduce risks to police officers and increase the efficiency and coverage of public safety operations.

The systems employed in this scenario have undergone technical certification and comply with existing national regulations governing public safety technology. Human operators remain nominally "in the loop" for decisions involving the use of force and arrest authorization. However, due to the volume of operations and the high level of confidence placed in automated systems, human oversight has increasingly become supervisory rather than deliberative. The human rights violations in this scenario arise not from technical malfunction, but from the expansion of state surveillance capabilities, over-reliance on automated risk assessments, and the gradual normalization of remote coercive policing.

Scenario

By 2029, several large metropolitan regions around the world have adopted integrated aerial public safety systems. Hundreds of autonomous flights occur daily, performed by networks of drones operated by municipal authorities, emergency services, infrastructure providers, and law enforcement agencies.

In the Republic of Astoria, a technologically advanced democratic state, the government has spent nearly a decade expanding its National Aerial Public Safety Network (NAPSN). Originally developed for emergency response, search and rescue operations, and disaster

management, the system has gradually evolved into a comprehensive platform for public surveillance and law enforcement.

Hundreds of autonomous surveillance drones patrol urban areas continuously. Equipped with high-resolution cameras, thermal imaging sensors, directional microphones, and behavioral analytics software, the drones collect and process vast amounts of information about public activities. Government officials maintain that the system has significantly reduced crime rates, improved emergency response times, and increased officer safety. Over time, however, the capabilities of the system have expanded.

Machine learning models are increasingly used to identify behaviors considered indicative of criminal activity, public disorder, or security threats. Individuals are assigned dynamic risk scores based on factors including their location history, known associations, prior interactions with law enforcement, participation in public demonstrations, and observed behavioral patterns. Although officials insist that the system does not make arrest decisions autonomously, recommendations generated by the AI are routinely relied upon by human operators.

At the same time, police departments have begun deploying remotely operated "intervention drones" capable of conducting arrests without requiring officers to be physically present. These systems employ loudspeakers, spotlight systems, non-lethal incapacitation devices, restraints, and remote communications equipment. Their use is justified primarily on the grounds of reducing risks to police officers during potentially dangerous encounters.

Initially, deployment of intervention drones was restricted to high-risk situations involving armed suspects. However, after several years without major incidents, their use gradually expanded to include routine arrests, public order enforcement, and crowd control operations.

On a Saturday afternoon in the summer of 2029, thousands of people gather in the capital city to protest recently enacted government surveillance legislation. The demonstration is lawful and peaceful. Throughout the event, participants are continuously monitored by aerial surveillance platforms operating above the city.

Among the demonstrators is Maya, a thirty-two-year-old community organizer and candidate for political office who has written extensively about government surveillance technologies. During the demonstration, she delivers a short speech criticizing the expansion of automated policing powers.

Several hours later, after the protest has concluded, an automated behavioral analysis system flags Maya as a person of interest. The system identifies her as a recurring participant in multiple political demonstrations, associates her with several activist organizations, and notes what it characterizes as "escalating anti-government rhetoric" in her public communications.

At the same time, another participant in the protest is arrested after minor acts of vandalism occur nearby. Through social network analysis and geolocation data, the system identifies Maya

as having been physically present within several hundred meters of the suspect during the protest. The system consequently elevates her risk score and recommends further investigation.

Because of increasing workloads and confidence in the system's predictive capabilities, the human analyst reviewing the recommendation authorizes surveillance continuation without conducting an extensive review.

Over the next forty-eight hours, autonomous surveillance drones monitor Maya's movements continuously. They observe her meeting colleagues, visiting her workplace, and attending a public lecture on civil liberties. Although no evidence of criminal activity is identified, algorithmic assessments continue to classify her as presenting an elevated risk of organizing future public disorder.

Three days later, authorities decide to detain several individuals for questioning regarding the vandalism investigation. Rather than dispatching police officers, the police operations center authorizes the deployment of intervention drones.

As Maya exits a commuter transit hub, two police drones descend nearby. Through loudspeakers, they announce that she is being detained and instruct her to remain still. Bystanders begin recording the encounter. Confused and frightened, Maya attempts to ask questions and steps backwards.

The automated threat assessment system interprets her movements as possible resistance. One drone deploys a non-lethal incapacitation device while another initiates a restraint procedure. Maya falls to the ground, suffering head injuries and a fractured wrist. She is subsequently transported to a detention facility, where licensed investigators quickly determine that she has no connection to the vandalism incident. Video footage of the arrest rapidly spreads across social media and international news networks.

A subsequent independent inquiry reveals that no individual component of the system malfunctioned. The surveillance systems identified precisely the individuals they were designed to identify. The predictive analytics system generated risk assessments consistent with its training data. Human supervisors approved each stage of the process according to existing procedures. The intervention drones operated within their certified technical parameters.

However, the inquiry concludes that the cumulative effect of pervasive surveillance, predictive policing, and remote coercive enforcement had fundamentally altered the relationship between citizens and the state. The report further concludes that the system had enabled forms of political surveillance and coercion that would have been difficult to achieve through traditional policing methods, while simultaneously reducing the opportunities for human judgment, discretion, and accountability that historically constrained the exercise of state power.

Preventative Actions

This scenario illustrates that the principal human rights risks associated with autonomous policing systems may arise not from technical failure, but from the expansion, normalization, and institutionalization of surveillance and coercive capabilities. Technologies developed for legitimate public safety purposes can gradually devolve into mechanisms that enable disproportionate state control over public life.

One possible approach is to establish clear legal limitations on the use of aerial surveillance technologies for public order monitoring and political activity. Surveillance of protests, political organizations, journalists, and civil society groups could require heightened legal thresholds, judicial authorization (similar to a warrant), and independent oversight. The collection and retention of data relating to lawful political activity might be subject to strict limitations or outright prohibition. Domestic U.S. case law suggests that using “sense-enhancing technology” like infrared cameras would constitute a search that requires a warrant.¹

A second approach would involve restricting or prohibiting the use of autonomous or remotely operated systems to exercise coercive policing powers. Governments could determine that certain state functions—including arrest, detention, and the use of force—require direct human presence and accountability. Under this model, drones could support police operations through observation and communication but would not be authorized to physically restrain, incapacitate, or arrest individuals.

A third approach would focus on maintaining meaningful human oversight. Human operators should not merely supervise automated decisions but should exercise independent judgment, possess sufficient information to challenge algorithmic recommendations, and retain both the authority and responsibility to reject automated assessments. Appropriate staffing levels, training, and organizational cultures are essential if human oversight is to remain substantive rather than symbolic.

Independent accountability mechanisms may also be required. Public reporting, external auditing, judicial review, and human rights impact assessments could help identify whether surveillance and enforcement systems disproportionately affect political activists, journalists, minority groups, or other vulnerable populations. Individuals subjected to automated surveillance or enforcement actions should have access to effective remedies, including notification, appeal procedures, and compensation where appropriate.

¹ *Kyllo v. U.S.*, 533 U.S. 27, 34 (2001). See *Leaders of a Beautiful Struggle v. Baltimore Police Dept.*, 2 F.4th 330 (4th Cir. 2021) (holding that even some drone surveillance systems can require warrants).

Finally, policymakers should recognize that technological efficiency and human rights protection are not always aligned objectives. Systems designed to maximize surveillance coverage, predictive accuracy, or officer safety may simultaneously reduce privacy, autonomy, accountability, and democratic participation. Decisions regarding the deployment of coercive technologies should therefore be treated not merely as technical questions, but as fundamental constitutional and human rights questions.

Observations

The following should be considered:

- Human rights impact assessments prior to the deployment of public safety drone systems.
- Explicit legal limits on surveillance of political activity, journalism, and peaceful assembly.
- Restrictions or prohibitions on autonomous or remotely operated use of force and arrest powers.
- Independent judicial authorization requirements for persistent aerial surveillance.
- Meaningful human oversight requirements for predictive policing and enforcement decisions.
- Regular auditing of algorithmic bias, discrimination, and disparate impacts.
- Transparency requirements concerning surveillance practices, retention periods, and enforcement outcomes.
- Public reporting on the use of predictive analytics and remote enforcement technologies.
- Independent complaint and redress mechanisms for individuals affected by automated policing systems.
- International standards governing the use of autonomous systems for surveillance, policing, and the exercise of state coercive powers.