Regulating Frontier AI with Simpol

The problem: no government can move first without fear of falling behind

Every major government now weighing AI regulation faces the same trap: tighten safety and testing requirements and you risk your AI industry falling behind competitors who regulate less. This is not a failure of national will — it is a collective action problem, and it explains why AI safety regulation across the US, the EU, China and elsewhere is being calibrated as much by competitive anxiety as by an honest assessment of risk.

Simpol's answer: remove the reason to move last

Simpol exists precisely for problems of this type. Governments agree, in advance and only in principle, to implement an agreed package of measures simultaneously, on the same date, once all or sufficient nations have signed on. If participating nations move together, no nation loses by moving first, and none gains by moving last. This is the same logic already underpinning Simpol's work on climate policy — AI regulation is simply the newest, and currently most urgent, instance of the same problem.

What an AI package would need to include

Work is underway, ahead of any formal negotiation, on what a Simpol AI package could look like in practice. Whatever detailed form it eventually takes, it would need to reflect the same principles that already govern every Simpol policy package:

1. **Global co-operation.** AI safety, testing and transparency measures would be implemented by participating nations simultaneously, only once all or sufficient governments have agreed to do so — so no nation or trading bloc suffers a competitive disadvantage for regulating responsibly.

2. **Subsidiarity.** Only measures that a nation could not safely adopt alone would be included. Purely domestic AI policy remains a matter for national governments and is expressly excluded from any Simpol AI package.

3. **Give and take.** AI safety and transparency commitments would be linked to a shared contribution from frontier AI compute, so that nations and companies bearing the largest share of frontier AI development also bear the largest share of the cost of proving compliance. Participation is intended to carry a benefit for every country, not only smaller or developing ones: the central benefit is verified confidence that competitors face the same requirements — something no nation can obtain by acting alone.

4. **Equality.** Democratic and non-democratic nations could participate on an equal basis; the initiative would not require any nation to alter its system of government, only to agree to the same simultaneous measures.

5. **Democracy.** Citizens in democratic countries would be invited to help shape the policies their government takes into negotiations, and asked to approve the framework for those negotiations, by referendum or other means suited to their constitution.

6. **Complementarity.** The Initiative is designed to work alongside — not against — existing efforts including the OECD AI Principles, the G7 Hiroshima AI Process, the EU AI Act, and the UN Global Dialogue on AI Governance, giving those efforts a mechanism for simultaneous, mutually-assured adoption where voluntary coordination alone has so far proved insufficient.

Scope: this is about frontier AI, not AI in general

"AI" here refers to frontier, general-purpose systems at the leading edge of capability — not to AI in general. The specific compute threshold that would operationalise this scope is a matter for the negotiated package itself, to be agreed and revisited on a fixed schedule as capability and hardware efficiency evolve. Narrower, beneficial applications — such as medical diagnostics, scientific tools like AlphaFold, or everyday consumer software — are not the target of this work and should continue to be encouraged.

Where this stands today

Simpol is currently developing the technical detail behind a possible AI package, including a proposal to fund independent verification of frontier AI safety claims through a modest levy on frontier AI compute. That proposal, and the open questions it still needs to resolve, is being tested through expert review and a negotiation simulation exercise and will only firm up into a specific policy proposal once it has been through that process. The detail, and the questions still open, are set out in the accompanying Annex document. None of this changes what is being asked today: only in-principle support for the process itself — via the Pledge for elected representatives, already signed by many UK MPs and by MPs in some other countries or, if you're a citizen, by joining the Simpol campaign.

What this means for you

If you are an MP or parliamentary candidate

- **Already signed the Simpol Pledge?** No further action is needed. AI regulation, like tax cooperation and climate policy before it, falls within the scope of the Pledge you have already signed. As with any Simpol issue, no further action is needed today. If the process proceeds you would first be asked, as with a trade negotiating mandate, to approve the boundaries under discussion — then, separately and later, to confirm any final agreement itself. Both are votes still to come, not commitments made today.

- **Not yet signed?** AI regulation is a current, live example of exactly the kind of problem the Simpol Pledge exists to solve. Signing costs you nothing beyond in-principle support, does not commit you to any specific policy text, and can be cancelled at any time. You can view and sign the Pledge at: simpol.org/sign-up/sign-up-page/pledge-page

If you are a citizen not yet signed on:

- Joining costs you nothing and takes only a minute. By joining, you commit to give strong voting preference at all future national elections to politicians who have signed the Pledge — turning your vote into leverage for exactly the kind of coordinated action AI regulation needs. Join the campaign now at: simpol.org/sign-up/sign-up-page/sign-up-citizens