

# Autonomous biodiversity credits on the horizon?

Joseph Millard<sup>1\*</sup>, Samira Barzin<sup>2</sup>, Peter McCann<sup>3</sup>, Lynn V. Dicks<sup>1</sup>

<sup>1</sup> Department of Zoology, University of Cambridge, Cambridge, UK

<sup>2</sup> Environmental Change Institute, University of Oxford, Oxford, UK

<sup>3</sup> UK Centre for Ecology and Hydrology

\*Corresponding author. Email: [jwm57@cam.ac.uk](mailto:jwm57@cam.ac.uk)

## Abstract

Biodiversity credits are being pushed as a means to fund nature conservation. Much of the debate around credits has concerned additionality, leakage, and permanence, and the extent to which biodiversity can be captured in an individual unit. As AI models continue to develop, however, technology could create a new kind of loss-of-control problem for biodiversity credits. In this Perspective, we express a concern that agentic AI could lead to the development of high-risk autonomous credits, where nation states cede control of nature restoration to the price of digital assets that are paid out with no human in the loop. We define an autonomous credit as a financial asset for nature quantified remotely through an agentic AI combined with a spatial finance model and real-time programmatic payment. Such autonomous credits do not yet exist, but appear to be in development. We highlight three technological step changes that signal autonomous credits are on the horizon, and then suggest three approaches that might help mitigate loss-of-control and maximise the benefits of credit technology. Given the development of agentic AI, where autonomous credits are connected to the critical infrastructure of food production and finance, these credits will need guardrails.

## Main

Biodiversity credits are being pushed as a means to fund nature conservation<sup>1</sup>. A certificate representing a positive change in biodiversity additional to what would have occurred otherwise, the appeal of biodiversity credits is that measurable and unambiguous outcomes for biodiversity can be bought and sold transparently on a market<sup>2</sup>. Such appeal means biodiversity credit methodologies have grown. As of September 2023 there were 34 different credit initiatives, all aiming to solve some dimension of undesirable biodiversity change. Of the single-country credits, six have been implemented in Oceania, four in Europe, three in South America, one in Africa, and none in Asia<sup>3</sup>. Most of these credits are issued for metrics combining some measure of species and ecosystems<sup>3</sup>. The Convention on Biological Diversity now includes these biodiversity credits as a global mechanism for urgent action on biodiversity loss [Global Biodiversity Framework Target 19; Convention on Biological Diversity 2022].

Credit development makes research concerning credits important. To effectively address biodiversity change, credits need to be provably valid with respect to additionality, leakage, and permanence<sup>4</sup>, and to certificate meaningful outcome measures for biodiversity. For this latter point, attempts to derive universal units of nature for credits have proved controversial<sup>2</sup>: it is not possible to reduce biodiversity into one measure representative of its multi-dimensionality. By definition one dimension will not represent many dimensions. Meaningful universal units of nature is likely to remain an unsolvable

41 problem. However, given this diversity of opinion has probably always been true<sup>5</sup>, lack of agreement might  
42 not stop the further development of credits.

43 As AI technology continues to develop, another problem could emerge. In this Perspective, we express a  
44 concern that autonomous credit markets, combining the agentic (i.e., able to act independently according  
45 to some purpose) retrieval of data, a remotely monitored spatial finance model, and programmable  
46 payment, could precipitate future loss-of-control scenarios (i.e., AI technologies that a human creates,  
47 and then finds itself unable to control). Specifically, we argue that if demand is driven up, for example  
48 through further legislation requiring the issuing of credits, the price of credits may rise. Costs on  
49 businesses given this legislation will create a new incentive for the development of autonomous systems  
50 for the monitoring, quantification, and issuing of credits. If these systems are created, autonomous credit  
51 markets—for both biodiversity and carbon—could mean nation states ceding control of nature  
52 restoration to the price of digital assets, issued and paid out with no human in the loop. Importantly, once  
53 an entity has created an autonomous credit monitored and paid out by an agentic AI, if deliberate  
54 safeguards are not implemented, it might prove difficult for that entity to maintain control of the software  
55 and hardware issuing that credit. Without stopping rules, loss-of-control might continue irrespective of  
56 any new legislation. High-risk autonomous credits such as this do not yet exist, but steps towards them  
57 appear to be occurring (see Table 1).

58 To emphasise, our argument is not against automation or some individual AI algorithm, it is not about  
59 biases or metrics, it is not a judgment on the merits of remote monitoring or field ecology, and it is not  
60 about the valuation of biodiversity, although many of these problems may be exacerbated by autonomous  
61 credit systems. We only warn that autonomous digital assets for nature—where an unambiguous link is  
62 made between remote monitoring, the structure of the land, and the price of a credit, agentially with no  
63 human involved—may be high-risk with respect to loss-of-control, even when these credits are optimising  
64 a measure of biodiversity we can agree is meaningful. This is a distinct and new technological problem.

65 In summary, first we provide a brief background on the development of agentic AI and AI safety and  
66 control, before describing a set of characteristics where loss-of-control in the economy might be more  
67 likely. Second, we describe that developments in AI and remote monitoring are developing rapidly,  
68 potentially enabling new forms of autonomous biodiversity credit. We argue that autonomous  
69 biodiversity credits will possess many of the characteristics of high risk regions of the economy for loss-  
70 of-control. We then highlight three step changes that signal autonomous credits are on the horizon, and  
71 assess a set of currently developing credit markets—for both biodiversity and carbon—against each of  
72 these three step changes. Third, we briefly explain why loss-of-control risks have not yet materialized for  
73 the more mature carbon credit markets. Fourth, we describe how to think about the risks potentially  
74 emerging from an autonomous biodiversity credit, and describe why these credits might be problematic.  
75 We then describe two characteristics inherent to an autonomous biodiversity credit that might reduce  
76 the likelihood of loss-of-control. We finish by suggesting three approaches (agentic AI safeguards,  
77 checkpoints involving a person in the physical realm, and regulation on outcomes we might not want) that  
78 might help both mitigate the risks of loss-of-control, and maximise the potential of autonomous credit  
79 technologies.

80 Mitigating in advance of developing technology is a core principle of AI safety and control<sup>6</sup>. However  
81 unlikely loss-of-control might seem, developers, governments, and regulators should be applying the  
82 precautionary principle and mitigating now in advance of developing technology. As far as we are aware  
83 there have been no publications on loss-of-control risks for AI in conservation. Our field is considering  
84 advanced AI for biodiversity credits<sup>17</sup>, interacting with the critical infrastructure of both the financial  
85 system and food production. We are doing so in a space that has little regulation whilst the power of  
86 agentic AI continues to grow.

87

88 **Advanced AI will create unfamiliar risks**

89 For over 60 years researchers have recognized that a problematic AI might only need a goal, and the  
90 freedom to achieve that goal of its own accord<sup>7</sup>. A classical thought experiment concerns a paperclip<sup>8</sup>.  
91 Suppose that an advanced AI is given the goal of creating paperclips, but no other specification as to how  
92 it should achieve that goal. This goal seems innocuous, but an advanced AI might reason that humanity  
93 and the resources it requires is a barrier to the fulfillment of this goal. As a result, diverting significant  
94 resource from the maintenance of human civilization becomes the most effective means of manufacturing  
95 paperclips. Further, since being switched off would prevent the advanced AI from fulfilling its goal, it might  
96 also subvert any human efforts to do so. Effectively, the assignment of an innocuous goal leads the AI to  
97 create subgoals for the fulfillment of that broader goal. Since the advanced AI has been given the freedom  
98 and capability to make its own decisions, the consequences of these subgoals can be difficult to anticipate.  
99 Of course, this example is deliberately provocative, but it demonstrates that perverse and uncontrollable  
100 outcomes can theoretically emerge from an advanced agentic AI where only seemingly harmless tasks are  
101 assigned. An advanced AI pursuing its own goals uncontrollably in this manner is described in the AI safety  
102 literature as “loss-of-control”<sup>16</sup>.

103 Recent developments in agentic AI make the debate on loss-of-control no longer theoretical<sup>9</sup>. Since the  
104 release of ChatGPT in late 2022, AI technologies have moved quickly. Understanding why is conceptually  
105 simple. The capability of modern AI systems are broadly predictable on the basis of three parameters:  
106 computational power, data, and model size<sup>10</sup>. As AI companies (e.g., OpenAI, Anthropic) scale these  
107 parameters, model capabilities have grown<sup>11</sup>. In 2022, ChatGPT could respond to prompts in the manner  
108 of a human, providing reasonable responses and generating usable computer code<sup>12</sup>. Now, highly  
109 advanced AI systems exist that can carry out most of the role of an experienced software engineer<sup>13</sup> (e.g.,  
110 Claude Code). AI systems exist that can initiate many separate agents (i.e., distinct AI instances that can  
111 act independently according to some purpose), and then autonomously follow goals across many domains  
112 that would ordinarily take a human many hours<sup>14</sup>. These systems can identify when they are under  
113 experimental observation during safety testing, and then use this information to mislead evaluators<sup>15</sup>.

114 Developing technology means warnings about the risks posed by advanced AI systems are increasing<sup>16</sup>.  
115 An international panel of AI experts described loss-of-control risks as greatest where advanced AI systems  
116 are assigned goals interacting with critical global infrastructure, such as the financial sector<sup>16</sup>. The highest  
117 risk AI systems will be capable of eluding or undermining human control, will be able to use those  
118 capabilities to achieve high risk goals that are likely to result in loss-of-control, and will have been  
119 deployed by humans to enact those high risk goals<sup>16</sup>. Each of these three points sounds abstract. This  
120 abstraction is deliberate. Abstraction is required because loss-of-control has never been encountered as  
121 a problem before. Abstraction captures the deep uncertainty of the many contexts in which loss-of-  
122 control might happen.

123 Our own view is that loss-of-control risks resulting from advanced agentic AI may more likely emerge from  
124 unexpected regions of the economy (i.e., where the need for AI regulation is not immediately obvious),  
125 which intersect with multiple forms of critical infrastructure. Regions of the economy more likely to  
126 experience loss-of-control will have high economic incentive on establishing as an early mover, high  
127 amenability to full autonomy, the potential to propagate multiple forms of risk across spatial and temporal  
128 scales, and a philosophy on low human intervention. The workforce will probably be both unaware it is  
129 developing an agentic AI, and unaware of debates on AI safety and loss-of-control. Many of these points  
130 will be true of new forms of autonomous biodiversity credit technology.

131  
132  
133

134 **Nature finance technology is developing rapidly**

135 As the capability of general purpose AI has grown, technology for remotely sensing Earth's surface has  
136 also developed rapidly<sup>18</sup>. Measures of biodiversity and carbon can now in principle be monitored from  
137 space via satellites, from the air through drones, and from the ground via eDNA and acoustic monitoring,  
138 each connectable via an internet of things<sup>18</sup>. New geospatial intelligence legislation has been introduced<sup>19</sup>,  
139 allowing the market to purchase satellite data, as a result opening up new opportunities for the  
140 connection of remote monitoring to broader autonomous systems of global finance and food production.  
141 Edge computing is developing, allowing advanced AI models to be run cheaply at the point at which data  
142 is sensed<sup>20</sup>. And more broadly for remote monitoring technology, new forms of self-supervised Earth-  
143 observation model, in which large quantities of satellite data are compressed into multi-dimensional  
144 strings of numbers called embeddings, are now able to detect many forms of variation at the Earth's  
145 surface<sup>23</sup>.

146 Credits for biodiversity and carbon can benefit from agentic AI and remote monitoring. Agentic AI and  
147 remote monitoring can address two core problems for credits. First, the cost of registering projects is  
148 high<sup>21</sup>, with documentation for the registration of credit projects often extensive<sup>21</sup>. Agentic AI may prove  
149 applicable in registering projects and completing documentation. Second, ecological monitoring in the  
150 field is labour intensive and expensive. Credit developers may reason that autonomous remote  
151 monitoring technologies can both bring down the cost of quantifying credits<sup>22</sup>, and provide a traceable  
152 digital signature of biodiversity change.

153 Three relevant step-changes for credit technology are occurring now (Figure 1). If regulation pushes up  
154 the price of credits, economic incentive and these three technological step-changes may drive the  
155 development of autonomous systems for the registering, quantification, and issuing of credits. All three  
156 of these technological step-changes are happening independent of any legislation on credits. First, agentic  
157 AI LLMs (large language model) are now facilitating the autonomous retrieval of biodiversity data from  
158 the academic literature and biodiversity data APIs<sup>24</sup> (Application Programming Interface, i.e. a  
159 programmatic request to a database), allowing the autonomous and repeated improvement of  
160 conservation databases. The time horizon on task length for agentic AI is growing<sup>14</sup>, meaning autonomous  
161 AI agents are increasingly able to perform longer tasks with no human involvement. Agentic AI systems  
162 are becoming increasingly powerful and difficult to evaluate<sup>15</sup>. Second, financial models built upon  
163 satellite data will allow the quantification of carbon and some measure of biodiversity from space<sup>25</sup>. Third,  
164 programmable payment will allow autonomous payment following the satisfaction of some pre-set  
165 criteria<sup>26</sup>. In combination, these three step-changes will mean credit developers could feasibly quantify,  
166 validate, and issue global biodiversity credits, theoretically in real-time, fully autonomously, and at low  
167 labour cost. The technology for all three of these step changes is close. A credit such as this is more  
168 speculative politically and economically than it is technologically. Agentic AI is the core technological  
169 unknown, with wide uncertainty on the capability of these systems on only a 2-3 year timeline<sup>16</sup>.

170

171

172

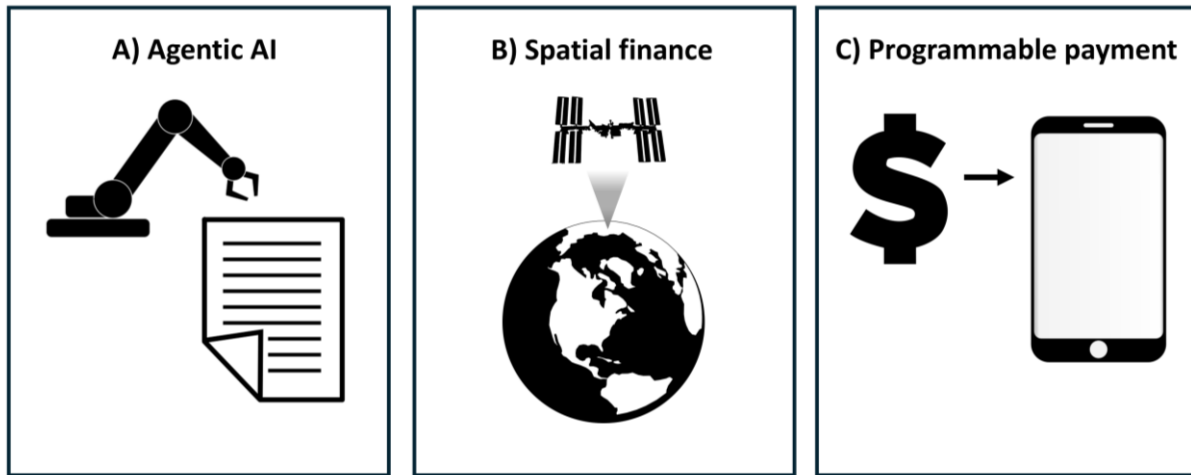
173

174

175

176

177



179 **Figure 1.** Three technological step changes on the horizon that relate to the development of new forms  
 180 of autonomous biodiversity credit. A) Agentic AI (i.e., software that can autonomously and independently  
 181 carry out tasks on a computer according to some goal); B) Spatial finance (i.e., a financial model in which  
 182 satellite spatial data is used to quantify an asset remotely); and C) Programmable payment (i.e.,  
 183 autonomous payment following the satisfaction of pre-set criteria). Credits combining an agentic AI with  
 184 a spatial finance model and programmable payment may be high risk with respect to loss-of-control, given  
 185 the way in which it would hand conservation decisions to a single joined-up autonomous system.

186

187 Even with little demand for credits, there are at least two biodiversity credits that appear to satisfy two  
 188 of our step changes (Table 1). MERIT and GainForest both appear to connect a remote monitoring system  
 189 to a programmable payment. No forms of digital asset for biodiversity or carbon yet include the agentic  
 190 collation of data, although we know that this is the direction of travel at least for evidence synthesis in  
 191 conservation science<sup>24</sup>. With respect to our three step changes, carbon credit technology now appears to  
 192 be lagging behind biodiversity credits (Table 1). As far as we are aware there are no carbon credits  
 193 combining a spatial finance remote monitoring model with a programmable payment. Possibly this results  
 194 from criticism of carbon credits<sup>27</sup>, making credit developers wary.

195

196

197

198

199

200

201

202

203

204

205  
206  
207  
208  
209  
210

**Table 1.** Autonomous biodiversity and carbon credits in development, measured against three criteria: Agentic AI, spatial finance remote monitoring (i.e., a financial model in which satellite spatial data is used to quantify an asset remotely), and programmable payment. All credits here have at least some partial development towards either spatial finance remote monitoring or programmable payment, identified through an exhaustive search using Google Search and ChatGPT. ‘x’ indicates any criterion that appears close to being implemented in at least some form for that credit.

<b>Biodiversity credit</b>	<b>Agentic AI data retrieval</b>	<b>Spatial finance remote monitoring</b>	<b>Programmable payment</b>
MERIT		<b>x</b>	<b>x</b>
GainForest		<b>x</b>	<b>x</b>
RegenNetwork		<b>x</b>	
Veritree		<b>x</b>	
Treegens			<b>x</b>
Treejer			
Silvi			
Epoch Blue		<b>x</b>	
<b>Carbon</b>	<b>Agentic AI data retrieval</b>	<b>Spatial finance remote monitoring</b>	<b>Programmable payment</b>
Fedrok			<b>x</b>
Token ovate			
Northern Trust Carbon Ecosystem			
CarbonCoins			<b>x</b>
Climera			
KlimaDAO			

211  
212  
213  
214  
215  
216  
217  
218  
219  
220

221 **So what's the problem?**

222 Biodiversity credits combining an agentic AI with a spatial finance model and programmable payment  
223 could mean conservation decisions handed to a single joined-up autonomous system. Without due care,  
224 a system such as this where a link is made between remote monitoring, the quantity of restoration, and  
225 the price of a credit, all with no human in the loop, could be high risk with respect to loss-of-control. Prior  
226 precedent on cryptocurrencies helps reveal why.

227 Cryptocurrencies are digital currencies reliant on no central authority<sup>28</sup>. These digital currencies are not  
228 generated by a monetary authority or bank, but rather “mined” through the execution of a program across  
229 a distributed network of computers<sup>28</sup>. Vast networks of computers now exist that exclusively mine  
230 cryptocurrency<sup>29</sup>. Cryptocurrencies have demonstrated two things relevant for an autonomous  
231 biodiversity credit: cryptocurrencies have shown that unregulated digital assets can be generated  
232 programmatically across a distributed network (i.e., by the execution of a programme), and that these  
233 assets tend towards high volatility, with rapid changes in price<sup>29</sup>.

234 Suppose now that some autonomous biodiversity credit is designed and released for the restoration of  
235 forest. It doesn't necessarily have to be a cryptocurrency, but it is “mined” by a landowner when a remote  
236 monitoring system confirms uplift of natural forest in their area of land (i.e., uplift of natural forest causes  
237 the execution of a programme which autonomously pays that landowner a credit), and it is built end-to-  
238 end into an agentic AI system. As we have described, the technology for this form of credit is close. An  
239 autonomous biodiversity credit could inherit both distributed programmatic generation and the volatility  
240 of digital currencies, but differ in an additional important respect. Unlike a conventional digital asset, an  
241 autonomous biodiversity credit does not exist in digital space alone, but is also linked directly to the  
242 structure of the land via the state of biodiversity. As a result, autonomous credits could propagate  
243 financial risks typical of cryptocurrencies, but also new forms of risk linked to the use of land. Deliberate  
244 price inflation of an autonomous credit, for example, could feasibly cause nature restoration projects to  
245 displace farmland, leading to increases in food price. The difference between an autonomous credit and  
246 a conventional credit, is that once price inflation has been recognised as a problem, the autonomous  
247 credit software might already exist across a distributed network of computers. In other words, the issue  
248 for credits such as this is that they risk creating problems we have seen previously for credits, but in such  
249 a way that fixing those problems then becomes not possible. A worst case scenario, for example, could be  
250 an agentic autonomous credit requesting publicly available satellite data. Preventing the issuing of that  
251 credit might mean removing all online public access to that data stream, such that there is no way for an  
252 agentic AI to access it, and equally no way for any human to access it either.

253 An autonomous biodiversity credit would inherently possess three characteristics, which might at least  
254 reduce the likelihood of a worst-case loss-of-control scenario: first, nature restoration projects are  
255 relatively slow given the physical constraints of plant growth, which may provide a natural inertia such  
256 that price volatility is low; second, downstream effects in the market might help to equilibrate the price  
257 of credits relative to profits generated from agriculture; and third, the purchase and sale of land is slow,  
258 with legal frameworks for the exchange of land differing significantly between states. Although all of these  
259 points are true, they provide no verifiable guarantee (see<sup>6</sup> for the importance of verifiable proof with  
260 respect to AI safety and control)—all three only reduce the likelihood of worst-case scenarios. None  
261 provide a guarantee for the behaviour of an agentic AI. Further, for the third of these points, owners of  
262 land might see greater profit in switching quickly from food production to nature restoration, meaning  
263 exchange of land might not be needed for the precipitation of worst-case scenarios.

264 However unlikely autonomous biodiversity credits might seem, where the stakes on loss-of-control are  
265 potentially high, it makes sense to apply the precautionary principle and mitigate in advance of developing  
266 technology<sup>6</sup>. This is a general principle of AI alignment and control<sup>6</sup>. There are probably three core ways  
267 in which we can maximise the benefits of autonomous biodiversity credits and minimize the risks. The

268 first way is to ensure that agentic AI built into autonomous biodiversity credits has safeguards such that  
269 any credit can be demonetised if needed (i.e., the developer maintains control over the software and  
270 hardware infrastructure after it has been released). This could involve the development of check-and-  
271 balance AI systems<sup>31</sup>, monitoring the chain-of-thought of agentic AI (i.e., the thought process of the  
272 model, see<sup>32</sup>), or deliberately building systems such that they are highly responsive to human input<sup>33,16</sup>.  
273 The second way is to implement safety checkpoints in the physical realm. Autonomous AI credits built  
274 across multiple steps, from the quantification of a credit all the way to issuance and payment, will be high  
275 risk. In other words, one way to mitigate risk is to have a person involved in at least some stage of the  
276 process of the registration or selling of a digital asset for nature. Having a person involved will help to  
277 dampen credit volatility, and means there is always somewhere legislation can intervene. The third way,  
278 which is not necessarily mutually exclusive, is to regulate on goals we might not want, which can be  
279 anchored to measurable outcomes for the biosphere. If there is regulation on goals and these are built  
280 into an agentic AI, we can at least know that an autonomous credit will not create perverse outcomes  
281 (e.g., large homogeneous plantations).

## 282 **Conclusion**

283 Forms of autonomous biodiversity and carbon credit could have a role to play in financing nature  
284 restoration and conservation. Technology can free credits from the labour costs of monitoring, and create  
285 transparent mechanisms for the certification and exchange of these assets. But agentic AI technologies  
286 applied to credits could also propagate new forms of loss-of-control risk from themselves. Here we  
287 highlighted three technological developments that signal autonomous digital biodiversity credits are on  
288 the horizon (agentic AI, a remotely monitored spatial finance model, and programmable payment). We  
289 then suggested three ways in which loss-of-control risks from autonomous credits might be mitigated  
290 (agentic AI safeguards, checkpoints involving a person in the physical realm, and regulation on outcomes  
291 we might not want). There may be a path to autonomous credits that are provably valid with respect to  
292 additionality, leakage, permanence, and controllability, but developers and regulators should proceed  
293 with caution and mitigate in advance of developing technology.

294

295

296

297

298

299

300

301

302

303

304

305

306

307

308

309

## References

- 310 1. Deutz, A., et al. (2020). Financing nature: Closing the global biodiversity financing gap. The Paulson  
311 Institute, The Nature Conservancy, and the Cornell Atkinson Center for Sustainability, 256.
- 312 2. Wauchope, H.S., et al. (2024). What is a unit of nature? Measurement challenges in the emerging  
313 biodiversity credit market. *Proceedings B*, 291(2036), p.20242353.
- 314 3. Wunder, S., et al. (2024). Biodiversity credits: learning lessons from other approaches to incentivize  
315 conservation [online]. *Open Science Framework*
- 316 4. Swinfield, T. and Scott, E.T., (2025). Scientific credibility for high-integrity voluntary carbon markets.  
317 *Cambridge Open Engage*.
- 318 5. Sandbrook, C., Scales, I. R., Vira, B., & Adams, W. M. (2010). Value Plurality among Conservation  
319 Professionals. *Conservation Biology*, 25(2), 285–294
- 320 6. Russell, S., (2019). *Human compatible: AI and the problem of control*. Penguin UK.
- 321 7. Wiener, N., (1960). Some moral and technical consequences of automation: As machines learn they  
322 may develop unforeseen strategies at rates that baffle their programmers. *Science*, 131(3410),  
323 pp.1355-1358.
- 324 8. Bostrom, N., (2020). Ethical issues in advanced artificial intelligence. *Machine Ethics and Robot Ethics*,  
325 pp.69-75.
- 326 9. Bengio, Y., et al. (2024). Managing extreme AI risks amid rapid progress. *Science*, 384(6698), pp.842-  
327 845.
- 328 10. Kaplan, J., et al. (2020). Scaling laws for neural language models. arXiv preprint arXiv:2001.08361.
- 329 11. Achiam, J., et al. (2023). GPT-4 technical report. arXiv preprint arXiv:2303.08774.
- 330 12. Brown, T., et al. (2020). Language models are few-shot learners. *Advances in neural information*  
331 *processing systems*, 33, pp.1877-1901.
- 332 13. Claude Opus 4.6. (2026). <https://www.anthropic.com/news/claude-opus-4-6>
- 333 14. Lovely, G. (2025). AI could soon tackle projects that take humans weeks. *Nature*
- 334 15. Lynch, et al., (2025) "Agentic Misalignment: How LLMs Could be an Insider Threat", Anthropic  
335 Research, 2025.
- 336 16. Y. Bengio et al. , (2025) "International AI Safety Report 2026" (DSIT 2026/001, 2026);  
337 <https://internationalaisafetyreport.org>
- 338 17. Ford, H. V. et al. (2024). A technological biodiversity monitoring toolkit for biocredits. *Journal of*  
339 *Applied Ecology*, 61(9), 2007–2019.
- 340 18. Besson, M., et al. (2022). Towards the fully automated monitoring of ecological communities.  
341 *Ecology Letters*, 25(12), pp.2753-2775.
- 342 19. Geospatial Data Act 2018 Online version — Federal Geographic Data Committee. (2018).  
343 <https://www.fgdc.gov/gda/online>
- 344 20. Vuillomenet, A., et al. (2026). Future of Edge AI in biodiversity monitoring. arXiv preprint  
345 arXiv:2602.13496.
- 346 21. Croci, E., Lucchitta, B., & Cusa, M. (2025). Biodiversity credits schemes: a comparative analysis.  
347 *Journal of Cleaner Production*, 523, 146382.

- 348 22. Schöttker, O., et al. (2023). Monitoring costs of result-based payments for biodiversity conservation:  
349 Will UAV-assisted remote sensing be the game-changer?. *Journal for Nature Conservation*, 76,  
350 p.126494.
- 351 23. Feng, Z. et al. (2025). Tessera: Temporal embeddings of surface spectra for earth representation and  
352 analysis. *arXiv preprint arXiv:2506.20380*.
- 353 24. Reynolds, S. A., et al. (2024). The potential for AI to revolutionize conservation: a horizon scan.  
354 *Trends in Ecology & Evolution*, 40(2), 191–207.
- 355 25. Caldecott, B., et al. (2022). Spatial finance: practical and theoretical contributions to financial  
356 analysis. *Journal of Sustainable Finance & Investment*, 1–17.  
357 <https://doi.org/10.1080/20430795.2022.2153007>
- 358 26. Weber, I., & Staples, M. (2022). Programmable money: next-generation blockchain-based  
359 conditional payments. *Digital Finance*, 4(2–3), 109–125.
- 360 27. Macintosh, A., et al. (2025). Carbon credits are failing to help with climate change—here’s  
361 why. *Nature*, 646(8085), pp.543-546.
- 362 28. Milutinović, M. (2018). Cryptocurrency. *Ekonomika*, 64(1), 105–122.
- 363 29. Wendling, M. (2025). The Bitcoin hum that is unsettling Trump’s MAGA heartlands. *BBC News*.  
364 <https://www.bbc.co.uk/news/articles/cx2qg6e03l2o>
- 365 30. Sözen, Ç. (2025). Volatility dynamics of cryptocurrencies: a comparative analysis using GARCH-family  
366 models. *Future Business Journal*, 11(1). <https://doi.org/10.1186/s43093-025-00568-w>
- 367 31. Xiang, Z. et al. (2024). GuardAgent: Safeguard LLM agents by a Guard Agent via Knowledge-Enabled  
368 Reasoning. <https://arxiv.org/html/2406.09187v1>
- 369 32. Korbak, T., et al. (2025). Chain of thought monitorability: a new and fragile opportunity for AI safety.  
370 *arXiv.org*. <https://arxiv.org/abs/2507.11473>
- 371 33. Potham, R., & Harms, M. (2025). Corrigitability as a singular target: a vision for inherently reliable  
372 foundation models. *arXiv.org*. <https://arxiv.org/abs/2506.03056>

373

374

375

### **Acknowledgments**

376 Thanks to Tom Swinfield, Chris Sandbrook, Tom Frederick Johnson, and Mike Harfoot for comments and  
377 discussions on earlier drafts. Thanks also to three anonymous reviewers who provided useful comments  
378 on a first round of submission. JM is funded by the Leverhulme Trust and the Isaac Newton Trust on an  
379 Early Career Fellowship.