The role of imagination in making water from moon rocks: How scientists use imagination to break constraints on imagination

MICHAEL T. STUART AND HANNAH SARGEANT

Abstract

Scientists recognize the necessity of imagination for solving tough problems. But how does the cognitive faculty responsible for daydreaming also help in solving scientific problems? Philosophers claim that imagination is informative only when it is constrained to be maximally realistic. However, using a case study from space science, we show that scientists use imagination intentionally to break reality-oriented constraints. To do this well, they first target low-confidence constraints, and then progressively higher-confidence constraints until a plausible solution is found. This paper exemplifies a new approach to epistemology of imagination that focuses on *sets* of imaginings (rather than individual imaginings), and *responsible* (rather than reliable) imaginings.

Keywords: imagination, space science, philosophy of imagination, philosophy of space science, internalism, externalism, reliabilism, responsibilism

1. Introduction

Philosophers interested in how we learn from imagination distinguish between 'transcendent' and 'instructive' uses of imagination (Kind and Kung 2016: 1). Transcendent uses of imagination, such as daydreaming, take us away from reality, while instructive uses, such as scientific thought experiments, are reality oriented. The 'puzzle of imagination use' asks how the same imagination can be put to both purposes (Kind and Kung 2016). A family of solutions begins from the idea that imagination can be used instructively only when its transcendent nature is constrained away. We argue against this view, claiming that transcendent imagination is often at the heart of instructive uses of imagination, for example when scientists build fictional versions of existing systems in their imaginations to investigate those very systems. The imagination is always constrained, of course, but it is also often *intentionally* in the business of breaking reality-oriented constraints, and therefore being transcendent. A good way to see this is by focusing not on individual uses of imagination but on sets of imaginings or processes of imagination.

Analysis Vol. XX | Number XX | XX 2024 | 1–14 doi: https://doi.org/10.1093/analys/anae015

This paper uses a case study to illustrate these ideas. Space scientists design instruments that do not yet exist, including probes, landers, orbiters and rovers, and all the sensors and experimental tools that each of these carries. They have launch windows that are years (or decades) in the future. Most of them are one-of-a-kind and the first-of-their-kind. In order to leave solid earth behind, space scientists must also leave behind the reality of their current problem contexts by using the imagination. Or so we claim.[1]

Scientists often decide whether an imagining was good or bad based on what it led to (Stuart 2022). An imagining that seems good at the time might have its status reversed if things fall through later. We examine a case based on the design and creation of an instrument that will fly on a future mission, because with such a case we can examine how scientists evaluate imaginings before they know the consequences. This helps us to identify norms that scientists use to decide how to imagine in the moment.

For our purposes, acts of imagination are cognitive actions that explore and manipulate aspects of a problem space, which are at least partially *imagistic* (i.e. more like perception than pure abstract thought), *creative* (i.e. they try to 'envision something new') and *freely variable* (i.e. not totally free but freer than most other kinds of mental action) (Sheredos and Bechtel 2020).

2. *ProSPA and the case of the disappearing water*

The ProSPA instrument takes its heritage from Open University (OU) spaceflight instruments, such as Ptolemy, which was flown onboard the Philae lander to comet 67P (Wright et al. 2007), and the Gas Analysis Package (GAP), which was flown onboard the Beagle 2 lander to Mars (Pullan et al. 2004). Both the Ptolemy and GAP designs include a carousel of ovens in which planetary surface samples are deposited and heated to detect any atmospheric gases and/or volatiles liberated from the samples. Chemical processing can also be used to produce what are known as analyte gases that are formed through reactions between the collected samples and onboard reference gases. The instruments utilize mass spectrometers to identify the gases, and reference gases are used to calibrate the instruments in flight.

With Beagle 2 on Mars in 2003, and Ptolemy having been launched towards comet 67P in 2004, the early 2000s saw focus returning to the Moon. The European Space Agency (ESA) was proposing a lunar lander concept to perform human exploration preparatory science, with a particular interest in the

---

[1] Here are a few other reasons why philosophers should pay attention to space science. (1) Space science is among the most interdisciplinary sciences, featuring geologists, chemists, (astro)biologists, (astro)physicists, astronomers, cosmologists and engineers. (2) Space science has special aims, including exploration. (3) Space science involves engineering in central and interesting ways. (4) Space science is of special existential relevance as it helps to define what (and where) humans are.

lunar poles and the availability of water (Carpenter et al. 2012). The search for lunar water (and its constituent oxygen) is of interest to the lunar science and exploration community as it could lead to the production of rocket propellant and supply some of the life support provisions required for a crewed lunar base (Lewis et al. 1993). Understanding the composition of this water will also help to determine how water came to be on the Moon.

Evidence for water on the Moon has recently grown (Colaprete et al. 2010, Mitrofanov et al. 2010, Meng et al. 2011, Li et al. 2018). The LCROSS impactor provided the first direct evidence of water on the Moon, when a spent rocket stage from the launch of a satellite was directed into Cabeus crater. Water was detected in the impact plume by a shepherding spacecraft, supporting the remote sensing evidence for water in polar craters.

An ESA package including a drill was initially proposed to ROSCOSMOS for inclusion in one of their upcoming lunar lander missions to search for lunar water. The payload has since been re-assigned to a NASA-led Commercial Lunar Payload Services (CLPS) mission. The ESA payload is known as the Package for Resource Observation and in Situ Prospecting for Exploration, Commercial exploitation and Transportation (PROSPECT) (Carpenter et al. 2014). PROSPECT was initially designed to identify water and other volatiles present in the regolith using two instrument packages. The PROSPECT Sample Excavation and Extraction Drill (ProSEED) is designed to take samples from beneath the lunar surface and deposit them into an analysis instrument called the PROSPECT Sample Processing and Analysis Suite (ProSPA). ProSPA is used to heat the samples and detect any released volatiles. The PROSPECT package is expected to be flown on the tenth CLPS mission in late 2026 (Figure 1).

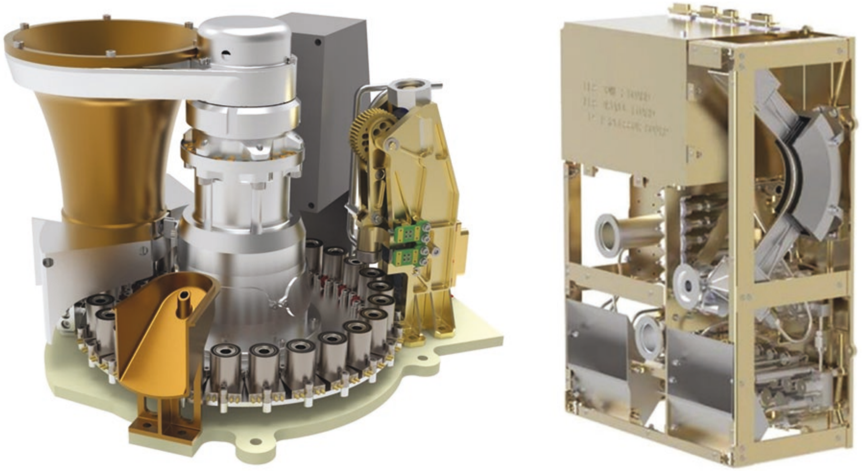

Figure 1 Rendering of the ProSPA payload with the Sample Inlet System (left) and Sample Analysis Suite (right). Credit: ESA.

An additional experiment was considered for ProSPA, which would extract oxygen from the regolith samples using the existing ProSPA design. Chemical extraction of oxygen from lunar regolith provides another source of oxygen other than water ice deposits. Some oxygen extraction methods require heating specific minerals (namely ilmenite) in the presence of reducing gases (Schlüter and Cowley 2020), and the OU lab theoretically had the hardware required to demonstrate that such experiments were feasible with the ProSPA design.

A PhD student joined the OU team in 2016 to investigate oxygen extraction with the ProSPA instrument. The team had backgrounds in mass spectrometry and lunar geology. Initially the project was not well defined. It began with a literature review of resource extraction techniques for use with lunar regolith. A trade-off study concluded that of the more than 20 techniques available to produce oxygen (Taylor and Carrier III 1993), hydrogen reduction was the most feasible, as ProSPA theoretically had the required hardware and hydrogen supply. At this point, nothing further than a thought experiment had been performed in considering this experiment for ProSPA.

The principle of the reaction is that hydrogen will bond with the oxygen produced by heating certain lunar minerals. However, this would be an equilibrium reaction, meaning that the water produced would need to be removed from the reaction site for the reaction to continue. This is generally achieved using a gas flowing system, that is, the hydrogen is flowed over the sample and the water is carried away (Keller et al. 2009, Kleinhenz et al. 2009, Lee et al. 2013). Owing to space constraints and the limitations of the existing design, ProSPA could not include a recirculating gas pump to create a gas flowing system. Instead an onboard 'cold finger' was proposed to condense the water away from the reaction site, enabling the reaction to continue (Figure 2).

An experimental setup was created to serve as an exploratory prototype of ProSPA, called a 'benchtop development model' (BDM). This replicated some of the key aspects of ProSPA's design (i.e. furnace, cold finger, pressure sensors and gas supply) (Figure 3). The purpose of the BDM was to perform example experiments relating to all the science goals, not just oxygen extraction. Experimental work is notoriously time-consuming and problem-laden, therefore a 'back-of-the-envelope' study was performed to check if the cold finger approach could work for the oxygen extraction experiments. This comprised a relatively simple set of calculations to determine how quickly gases would move from one end of a pipe to another when applying relevant temperatures and pressures. The timescale mattered, because power limitations on the Moon meant that ProSPA could run the reaction for up to four hours only. Thus, if the BDM took longer than four hours for water vapour to migrate through the system, the reaction would be deemed unsuccessful. The back-of-the-envelope calculation gave promising results, which justified preliminary experimental work.

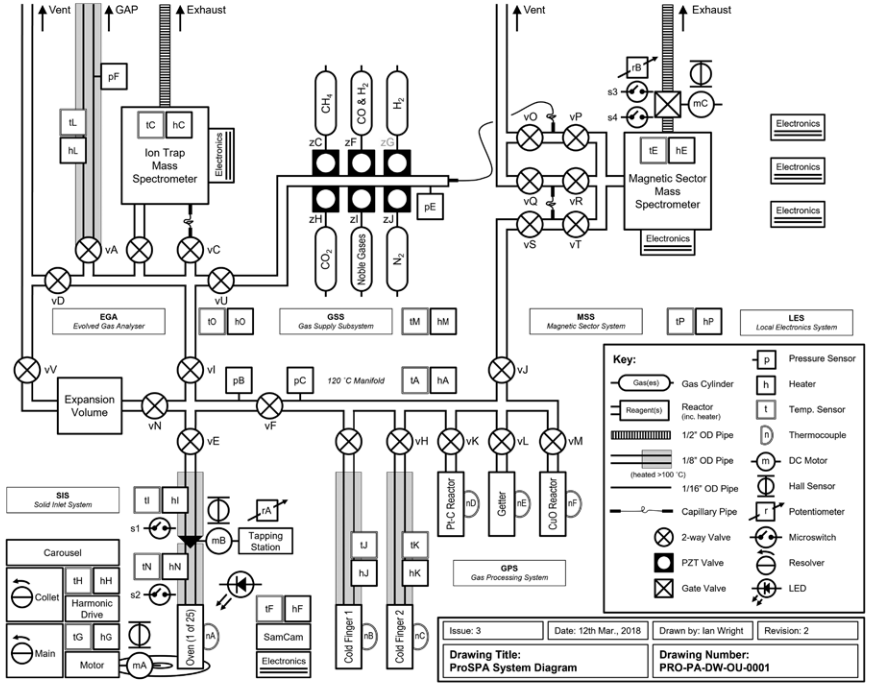

Figure 2 ProSPA's design at the time of writing. ProSPA is outfitted with a range of onboard gases (H_2, CO, CO_2, CH_4, N_2 and noble gases), a gas control system, pressure sensors, two cold fingers, two mass spectrometers and furnaces capable of heating samples to 1,000°C (Sargeant 2020, 28).

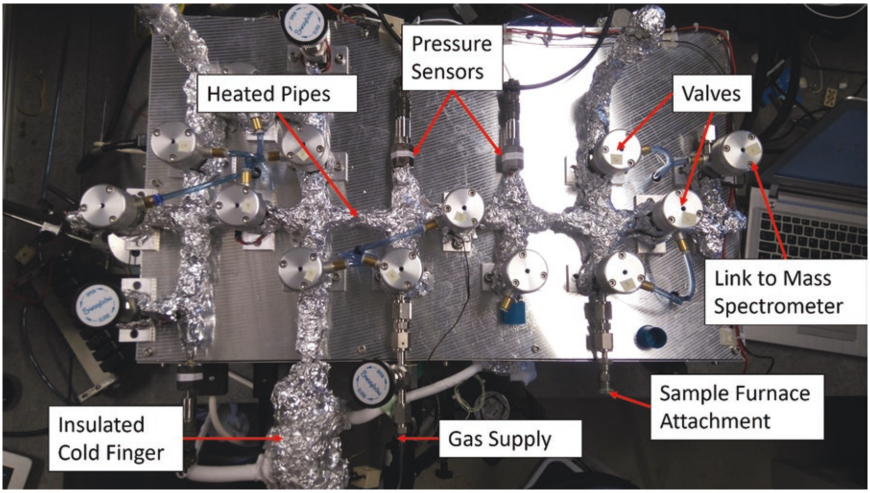

Figure 3 Top view of the BDM experimental setup. Heater tape and insulation were added to the pipework after initial experiments indicated that water was condensing in other places than the cold finger (modified from Sargeant et al. 2020a).

After calibration of the BDM, some ilmenite (the mineral of interest for this reaction) was heated in the furnace and combined with hydrogen with the aim of collecting water at the cold finger. To check whether water was collected, at the end of the reaction the cold finger was heated up and the pressure was recorded. A rise in pressure meant water had been collected. However, there was an unexpected result: after a rise in pressure, there was an immediate drop. The more ilmenite the team used in the reaction, the higher the spike in pressure from the release of water from the cold finger, always followed by a drop. At this point the scientists began imagining. One idea was that there might be a gas leak, allowing water to leak out of the system. This was perhaps the most obvious explanation, but it was proven wrong through leak testing.

When no leak was found, the scientists began imagining the different processes that they expected to happen in the system, one by one. It was assumed that water was in the vapour phase unless it was at the cold finger. However, upon imagining the setup in more detail, the team realized that water would not necessarily be in the vapour phase at the temperatures and pressures considered because water could also be condensing on other parts of the pipework, which would prevent it from being recorded by the pressure sensors. A solution was developed in imagination: if some heater wire was wrapped around the pipework (not including specific temperature-sensitive components), perhaps the water could be maintained as vapour everywhere but the cold finger. This was done, the experiments were repeated and the results were more promising.

However, the measured pressure was still not as expected. As before, the team imagined the entire process, step by step, trying to find (in imagination) other unheated parts of the system that might still act like water traps. At this point, the team imagined a new idea, which required re-thinking the entire instrument. Rather than heating individual components, they considered heating *everything* to a uniform temperature. That meant sourcing materials, sensors and valves that could all be heated. Instead of putting an oven in the instrument, the instrument would now go inside an oven. The new in situ resource utilization system (ISRU-BDM) was built inside a box made from insulating board and heated with heater elements used to heat a conventional kitchen oven (shown in Figure 4). Again this led to improvements in the measurements of the amount of water produced.

However, the amount of water produced was *still* not the amount they should have been seeing. Again the entire process was re-imagined, mostly visually, going through each part of the process, thinking in terms of molecules of water, oxygen and hydrogen (usually pictured as little balls bouncing around) and their reactions. The amount of water produced was always less than the amount of hydrogen used. Where was the hydrogen going if it was not combining to form water? A final breakthrough came via imagination. The more ilmenite they used in the reaction, the more water was 'lost'.

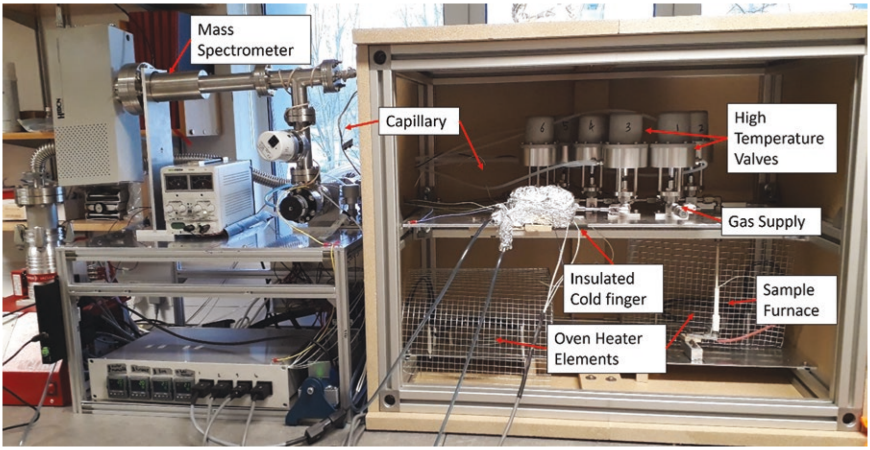

Figure 4. Front view of the ISRU-BDM with the oven door removed (Sargeant et al. 2020b).

Thinking very carefully about this difference was the key to overcoming a certain dogma in their previous imaginings. Previously the experiments were treated as if they were finished within the set time frame, but the reactions were not *finished*; they were artificially *halted* by the four-hour operating limit. What happens when the reaction is halted? Focusing on this moment in imagination made it possible to see that when the reaction was halted, the water that was then being produced would be trapped amid the grains of the sample. That final amount of water would not be able to reach the cold finger to be ‘counted’ in the final yield. When the samples had more ilmenite, there would be more residual water trapped in the sample that was just being formed, and thus missed from the final total (Sargeant et al. 2021). Quantifying these effects finally resulted in more accurate estimates of yield for the final operation of the flight instrument.

One way to understand the practice of space instrumentation is as solving a series of theoretical and technical problems. Many of these problems will require imagination to solve, and many of them are solved through collaborative acts of imagination, or imagination distributed across humans, machines and material instruments.

3. *When is an imagining in space science ‘good’?*

One way to evaluate an imagining is to look at its consequences (Stuart 2022). If the imagining led to a desired result, it was good. Because this appears to be the dominant way that scientists evaluate imaginings, it is difficult if not impossible to provide a set of rules that determine for any imagining whether it is good or not. After all, an imagining might produce good consequences by breaking our best rules (Stuart 2020).

However, there are different senses of 'good' we can distinguish. One important distinction is between *reliable* imaginings (which we can define as those that reliably produce good consequences) and *responsible* imaginings (which we can define as those that are rational from the perspective of the agent or group, perhaps because they are expected to produce good consequences). Scientists often do not know the consequences of a course of imagining, because they cannot foresee whether it will solve a given problem. However, scientists can still imagine responsibly, in the sense that they can imagine in a way that can be expected to have good consequences. Making and communicating judgements about which imaginings are responsible is important for signalling that an imagining should be developed further, to encourage more imaginings along similar lines, and to provide positive reinforcement.

How can scientists determine whether they are imagining responsibly? Translating the dominant position in epistemology of imagination into the above framework, philosophers claim that we imagine responsibly when we constrain away the transcendent aspect of imagination (see e.g. Nersessian 1992, 2007, Byrne 2005, Miščević 2007, Gregory 2010, Currie 2016, Kung 2016, Williamson 2016, Berto 2017, Kind 2018, 2021, Lam 2018, Canavotto et al. 2020; for critical discussion, see Stuart 2020). For example, Kind draws an analogy with computer simulations, which are used responsibly when their inputs contain accurate representations of the target system, and their structure manipulates those inputs in a way that accurately reflects how the target system is structured. Likewise, we imagine responsibly when we constrain our imaginings by using background knowledge, by employing only accurate representations of the target system and by unfolding the imagining as the real system would unfold. Kind writes 'beliefs about the world infuse my imaginings. In doing so, they act as constraints on my imagination, just as pre-programmed variables set constraints on computer simulations. When I set myself these imaginative projects, I don't take myself to be completely free. In fact, I don't take myself to be free at all' (2018: 243). The freedom of transcendent imagination is what makes the imagination unreliable. To imagine responsibly, that freedom must be eliminated.

Even supposing that the imagination could be fully constrained and that doing so might have epistemic benefits, we should note that many instructive uses of imagination are transcendent, and it is epistemically good that they are.[2] Of course, scientists do not necessarily *want* to break constraints: it is much easier when background knowledge straightforwardly implies a solution to a problem and the usual methods work. However, this is not always what happens.

---

[2] For what it is worth, most computer simulations are also transcendent. Scientists need freedom to depart from what they believe to be true about the target system to make useful computational models, at least because computer programs all have hardware and software limitations that make it difficult to model real-world (especially analogue) systems accurately.

As the case of ProSPA shows, scientists are sometimes forced to go beyond the stated aims and methods of the funder's commission, reject existing methods, reject assumptions given by experience and background theoretical knowledge, and rethink the entire instrument itself. Transcendent imagination is the source of our freedom to try something new when we cannot adhere to all the relevant constraints. The norm for imagining well in space science cannot be: satisfy *all* the relevant constraints that could apply. There must be some other norm, one that demands the breaking of constraints in a responsible way.

There are different ways to go about choosing which constraints should be broken and which should be obeyed. How do space scientists do it? In response to each of the problems outlined above, space scientists re-imagined their system, usually more than once. On their first try, they broke the constraints that they had the least confidence in. For example, they believed there was no leak in the system, but they were not very confident about that, so they imagined what they would find if there was a leak. If that suggested a plausible solution, they stopped imagining, and started calculating, modelling or experimenting. Once that possibility was cut off, they would re-imagine the system again, this time breaking a different constraint, or the same constraint as before, but in a different way or by a bigger margin. Typically only one constraint, or a set of similar constraints, will be broken at a time. Constraints in which scientists have a high degree of confidence, including strongly supported theoretical generalizations, are broken only as a last resort. Doing so will be entertained only when all other constraints have been explored.

This method for imagining responsibly assumes that low-confidence constraints are better to break first, in the epistemic sense of 'better'. Why? It is reasonable to mistrust a constraint when it is based on, for example, less evidence compared with one supported by dozens or hundreds of empirical studies. Of course, in rare cases, it will indeed be a high-confidence constraint that is incorrect and needs to be broken. But it is not rational from the perspective of the scientist to break the best-established constraints first when a problem arises. More likely, the scientist made a mistake in their own thinking that they did not notice.

The norm we have outlined is not sufficient to tell a scientist exactly which constraints to break and in what order, for two reasons. First, it will not always be clear, even to the scientist, which constraints they have the least confidence in. Much imagination is unconscious, at least in the sense that it is inaccessible to introspection (Stuart 2019a). Many constraints operating on unconscious processes will also not be accessible to introspection, and so it will be difficult to rank them in terms of confidence. Second, even with a complete confidence ranking, it is not clear how to compare confidence in complex cases, at least because the reasons for being confident about one constraint might be of a very different kind from the reasons we have for another. For example, how do you compare confidence grounded in personal ability to confidence grounded in peer-reviewed literature? Complexity of this kind can make it difficult to decide whether to break a single high-confidence constraint or several low-confidence constraints, for example, or whether to break two constraints by a little bit each, or one by a lot, etc.

In sum, scientists must choose not only which constraints to break, but how many to break, and in what order, and by how much to break each constraint. And they must do this with only partial information. This is a complicated minimization problem. At least one constraint must be broken, and in principle any constraint can be broken, but as few as possible should be broken, and those that are broken should be broken as little as possible; and it will not be clear until the imagining has finished which constraints were the right ones to break.

We think this norm is descriptively adequate in that scientists do appear to imagine in line with it. Now, we want to discuss some epistemological reasons for and against it. First, we should say that it is ethically and epistemically good that this norm recognizes the importance of imagination for breaking constraints. Many scientists feel that imagination is not, or should not be, part of their work, and an explicit norm that celebrates its importance could be very helpful (Stuart and Sargeant forthcoming). It is also good that this norm recognizes the uncertainty inherent in deciding which constraints to break, and how to break them. This emphasizes the freedom of imagination in science, and allows room for productive mistakes. Under conditions of perfect knowledge, perhaps we could determine in advance which situations should be imagined and in which order, enabling the mechanization of scientific imagination. While it is inevitable that more imagination will be outsourced to machines (Chandrasekharan et al. 2013, Shinod 2021, Stuart 2023), we can still afford to celebrate and defend imagination as one of the most human aspects of science (Stuart 2021).

One negative consequence might be that this norm promotes a kind of conservativeness that is undesirable for epistemological and ethical reasons (Currie 2019, Stuart 2019b, Stuart and Sargeant forthcoming). That is, it recommends breaking constraints as little as possible, and moving reluctantly from low-confidence to higher-confidence constraints. There are reasons to be conservative, and properly balancing conservativeness with open-minded exploration is difficult. But we must be very careful with conservativeness when applied to *imagination*, the one faculty capable of getting us out of the boxes that conservativeness puts us in.

Another point in favour of the norm is that while we cannot specify exactly how to satisfy it in advance, through the process of trying to adhere to it, scientists will often learn more about their own confidence levels through their imaginative acts of trial and error, and this is a useful epistemic side-effect of the process.

Lastly, this norm has enough bite to enable criticism of existing imaginative practices. For example, scientists might choose to break constraints for practical or aesthetic reasons rather than epistemic ones. A scientist might choose to break the constraint on representational accuracy, not because they have low confidence in that constraint in the present context, but because they would rather explore an epistemically worse but easier solution. That is not responsible imagining.

In sum, our inquiry yields a procedural ideal for responsible scientific imagining: when using imagination to find a solution, identify the relevant constraints as carefully as possible, and break them one at a time (or a few at a time), starting with those in which there is the least confidence. As each attempt fails, break the constraints more radically, and break different constraints, until a solution is found. This kind of advice is sorely needed for scientific imagination, which is a skill that scientists must learn, and there is currently no generally accepted procedure for learning how to use it (Stuart and Sargeant forthcoming).

4. *Conclusion*

Space science has features that should make it very interesting for philosophers of science interested in imagination. In this paper, we have identified a norm that enables space scientists to evaluate the responsibility of imaginings without knowledge of the consequences of those imaginings. As long as a scientist imagines according to the above-described norm, they are imagining responsibly, and therefore doing something praiseworthy. We do not claim that this norm extends to other fields of science, but we see no reason why it would not.

Two issues are worth flagging here. First, we differentiated between reliable and responsible imaginings. The concept of reliable imagining seems to be externalist, while the concept of responsible imagining seems to be internalist. If this is correct, perhaps resources from epistemology concerning that distinction would be useful here. Second, the practice of space science, like most modern science, is radically interdisciplinary, and how scientists imagine *together* must be taken into account to flesh out the above considerations fully.

In sum, the existence of the above-identified norm tells against any epistemology of imagination that requires constraining away imaginative freedom for an imagining to be instructive. Instead the ability of imagination to facilitate mental experimentation (and thereby surprise us; see French and Murphy 2023) must be counted as epistemically valuable, and its freedom to do this must be part of the story of what makes science epistemically productive. In a breakthrough, what is broken through are the constraints that define our starting point. If breakthroughs are epistemically good, and we assume they are, we ought to reject any epistemology of scientific imagination that rejects constraint-breaking.[3]

---

[3] Mike Stuart thanks audiences at the IPMC and the National Taiwan University and participants at conferences: CSHPS, SPSP (Ghent) and 'Imagination and its Constraints' (University of Parma).

Funding

Hannah Sargeant acknowledges funding from the Science and Technology Facilities Council (STFC) (#ST/N50421X/1). ProSPA is a programme of, and funded by, the European Space Agency (ESA).

University of York
UK
Aerospace Engineering
University of Leicester
UK
mike.stuart.post@gmail.com

References

Berto, F. 2017. Impossible worlds and the logic of imagination. *Erkenntnis* 82: 1277–97.

Berto, F. 2021. Taming the runabout imagination ticket. *Synthese* 198: 2029–43.

Byrne, R.M.J. 2005. *The Rational Imagination: How People Create Alternatives to Reality*. Cambridge, MA: MIT Press.

Canavotto, I., F. Berto and A. Giordani. 2020. Voluntary imagination: a fine-grained analysis. *Review of Symbolic Logic* 15: 1–26.

Carpenter, J.D., S. Barber, P. Cerroni, R. Fisackerly, A. Fumagalli, B. Houdou, C. Howe, P.G. Magnani, A. Morse, E. Monchieri, P. Reiss, L. Richter, F. Rizzi, S. Sheridan, L. Waugh and I.P. Wright. 2014. Accessing and assessing lunar resources with PROSPECT. *Annual Meeting of the Lunar Exploration Analysis Group*. <https://oro.open.ac.uk/41598/>

Carpenter, J.D., R. Fisackerly, D. de Rosa and B. Houdou. 2012. Scientific preparations for lunar exploration with the European Lunar Lander. *Planetary and Space Science*, 74: 208–23.

Chandrasekharan, S., N. Nersessian and V. Subramanian. 2013. Computational modeling: is this the end of thought experiments in science? In *Thought Experiments in Science, Philosophy and the Arts*, eds. M. Frappier, L. Meynell, and J.R. Brown, 239–60. London: Routledge.

Colaprete A., P. Schultz, J. Heldmann, D. Wooden, M. Shirley, K. Ennico, B. Hermalyn, W. Marshall, A. Ricco, R.C. Elphic, D. Goldstein, D. Summy, G.D. Bart, E. Asphaug, D. Korycansky, D. Landis and L. Sollitt. 2010. Detection of water in the LCROSS ejecta plume. *Science* 330: 463–8.

Currie, G. 2016. Imagination and learning. In *The Routledge Handbook of Philosophy of Imagination*, ed. A. Kind, 427–39. London: Routledge.

Currie, A. 2019. Existential risk, creativity & well-adapted science. *Studies in History and Philosophy of Science Part A* 76: 39–48.

French, S. and A. Murphy. 2023. The value of surprise in science. *Erkenntnis* 88: 1447–66.

Gregory, D. 2010. Conceivability and apparent possibility. In *Modality: Metaphysics, Logic, and Epistemology*, eds. B. Hale and A. Hoffmann, 319–46. Oxford: Oxford University Press.

Keller, B.W., D.L. Clark and J.A. Kirkland. 2009. Field test results of the PILOT hydrogen reduction reactor. *AIAA SPACE 2009 Conference & Exposition* 6475. <https://arc.aiaa.org/doi/abs/10.2514/6.2009-6475>

Kind, A. 2018. How imagination gives rise to knowledge. In *Perceptual Imagination and Perceptual Memory*, eds. F. Macpherson and F. Dorsch, 227–46. Oxford: Oxford University Press.

Kind, A. and P. Kung. 2016. Introduction: the puzzle of imaginative use. In *Knowledge Through Imagination*, eds. A. Kind and P. Kung, 1–38. Oxford: Oxford University Press.

Kleinhenz, J.E., Z. Yuan, K. Sacksteder and J. Caruso. 2009. Development of a reactor for the extraction of oxygen and volatiles from lunar regolith. *47th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. <https://doi.org/10.2514/6.2009-1203>

Kung, P. 2016. Imagination and modal epistemology. In *The Routledge Handbook of Philosophy of Imagination*, ed. A. Kind, 427–39. London: Routledge.

Lam, D. 2018. Is imagination too liberal for modal epistemology? *Synthese* 195: 2155–74.

Lee, K.A., O. Lara, A.J. Paz, R. Michael and M.S. Thomas. 2013. The ROxygen project: outpost-scale lunar oxygen production system development at Johnson Space Center. *Journal of Aerospace Engineering* 26: 67–73.

Lewis, J.S., D.S. McKay and B.C. Clark. 1993. Using resources from near-earth space. In *Resources of Near-Earth Space*, eds. J. S. Lewis, M.S. Matthews and M.L. Guerrieri, 3–14. Tucson, AZ: University of Arizona Press.

Li, S., P.G. Lucey, R.E. Milliken, P.O. Hayne, E. Fisher, J-P. Williams, D.M. Hurley and R.C. Elphic. 2018. Direct evidence of surface exposed water ice in the lunar polar regions. *Proceedings of the National Academy of Sciences* 115: 8907–12.

Meng, Z., C. Shengbo, L. Peng, W. Zijun, Y. Lian and Z. Chao. 2011. Research on the distribution and content of water ice in lunar pole regions using clementine UVVIS data. *Journal of Earth Science* 22: 595–600.

Miščević, N. 2007. Modelling intuitions and thought experiments. *Croatian Journal of Philosophy* 7: 181–214.

Mitrofanov, I.G., A.B. Sanin, W.V. Boynton, G. Chin, J.B. Garvin, D. Golovin, L.G. Evans, K. Harshman, A.S. Kozyrev, M.L. Litvak, A. Malakhov, E. Mazarico, T. Mcclanahan, G. Milikh, M. Mokrousov, G. Nandikotkur, G.A. Neumann, I. Nuzhdin, R. Sagdeev, V. Shevchenko, V. Shvetsov, D.E. Smith, R. Starr, V.I. Tretyakov, J. Trombka, D. Usikov, A. Varenikov, A. Vostrukhin and M.T. Zuber. 2010. Hydrogen mapping of the lunar south pole using the LRO neutron detector experiment LEND. *Science* 330: 483–86.

Nersessian, N.J. 1992. In the theoretician's laboratory: thought experimenting as mental modeling. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1992: 291–301.

Nersessian, N.J. 2007. Thought experiments as mental modelling: empiricism without logic. *Croatian Journal of Philosophy* 7: 125–61.

Pullan, D., M.R. Sims, I.P. Wright, C.T. Pillinger and R. Trautner. 2004. Beagle 2: the exobiological lander of Mars Express. In *Mars Express: The Scientific Payload*, ed. A. Wilson, 165–204. Noordwijk: European Space Agency.

Sargeant, H.M. 2020. *Water from Lunar Regolith: Reduction by hydrogen for a small-scale demonstration of in situ resource utilisation for the Moon*. PhD thesis, The Open University.

Sargeant, H.M., F.A.J. Abernethy, M. Anand, S.J. Barber, P. Landsberg, S. Sheridan, I.P. Wright and A.D. Morse. 2020a. Feasibility studies for hydrogen reduction of ilmenite in a static system for use as an ISRU demonstration on the lunar surface. *Planetary and Space Science* 180: 104759.

Sargeant, H.M., F.A.J. Abernethy, S.J. Barber, I.P. Wright, M. Anand, S. Sheridan and A.D. Morse. 2020b. Hydrogen reduction of ilmentite: towards an in situ resource utilization demonstration on the surface of the Moon. *Planetary and Space Science* 180: 104751.

Sargeant, H.M., S.J. Barber, M. Anand, F.A.J. Abernethy, S. Sheridan, I.P. Wright and A.D. Morse. 2021. Hydrogen reduction of lunar samples in a static system for a water production demonstration on the moon. *Planetary and Space Science* 205: 105287.

Schlüter, L. and C. Aidan. 2020. Review of techniques for in-situ oxygen extraction on the moon. *Planetary and Space Science* 181: 104753.

Sheredos, B, and W. Bechtel. 2020. Imagining mechanisms with diagrams. In *The Scientific Imagination: Philosophical and Psychological Perspectives*, eds. A. Levy and P. Godfrey-Smith, 178–209. Oxford: Oxford University Press.

Shinod, N.K. 2021. Why computer simulation cannot be an end of thought experimentation. *Journal for General Philosophy of Science* 52: 431–53.

Stuart, M.T. 2019a. Towards a dual process epistemology of imagination. *Synthese* 198: 1329–50.

Stuart, M.T. 2019b. Everyday scientific imagination: a qualitative study of the uses, norms, and pedagogy of imagination in science. *Science & Education* 28: 711–30.

Stuart, M.T. 2020. The productive anarchy of scientific imagination. *Philosophy of Science* 87: 968–78.

Stuart, M.T. 2021. Telling stories in science: Feyerabend and thought experiments. *HOPOS* 11: 262–81.

Stuart, M.T. 2022. Scientists are epistemic consequentialists about imagination. *Philosophy of Science* 90: 518–38.

Stuart, M.T. 2023. The future won't be pretty: the nature and value of ugly, AI-designed experiments. In *The Aesthetics of Scientific Experiments*, eds. M. Ivanova and A. Murphy, 215–33. London: Routledge.

Stuart, M.T. and H. Sargeant. forthcoming. Inclusivity in the education of scientific imagination. In *Building Inclusive Ethical Cultures in STEM*, eds. E. Hildt, K. Laas, C. Miller and E. Brey. London: Routledge.

Taylor, L.A. and W.D. Carrier III. 1993. Oxygen production on the moon: an overview and evaluation. In *Resources of Near-Earth Space*, eds. J.S. Lewis, M.S. Matthews and M.L. Guerrieri, 69–108. Tucson, AZ: University of Arizona Press.

Williamson, T. 2016. Knowing by imagining. In *Knowledge Through Imagination*, eds. A. Kind and P. Kung, 113–23. Oxford: Oxford University Press.

Wright, I.P., S.J. Barber, G.H. Morgan, A.D. Morse, S. Simon, D.J. Andrews et al. 2007. Ptolemy – an instrument to measure stable isotopic ratios of key volatiles on a cometary nucleus. *Space Science Reviews* 128: 363–81.