If Platonic abstract objects exist and if humans can process them, then we should be able to find experimental evidence for this at the Neuroscience/Psychology interface

Martin Korth*

IVV NWZ, WWU Münster, Wilhelm-Klemm-Str. 10, 48149 Münster (Germany)

* Corresponding author email: dgd@uni-muenster.de

v1, 2025/11/28 – on the occasion of the U Bamberg 2025 Intl. Workshop on Computational Ontology & Metaphysics

[This manuscript is supposed to end up as a book chapter; it does not yet have proper references.]

Over the last years, 'sub-symbolic' deep neural network (DNN) based artificial intelligence (AI) systems have run into all sorts of problems, especially with 'hallucinations' and in general with reliability.(REFs) One way to conceptualize this is to understand sub-symbolic AI to lack proper world models,(REFs) i.e. some type of proper mapping of what's supposedly real and what's not – though it might be more helpful to concede that such AI systems have 'fuzzy' implicit world models, that are superpositions of all proper world models compatible with the statistics of the language token use in the training data.

Understood in one of these ways, it is tempting to think that the problems of these models can be solved by adding proper world model information, only that the generation and operation of open, extended world models is exactly the task at which 'classical' symbolic AI approaches failed.(REFs) Combined sub-symbolic and symbolic approaches (like for instance 'Neuro-AI' or similar)(REFs) are thus most likely doomed to run not only into some of the problems of both sub-symbolic, as well as symbolic systems, but even into additional problems related to the interfacing of the two.

For this interfacing, no generally applicable approach is known yet and given that (even after at least decades of research) our knowledge of how humans generate or access (layers of) symbol systems and world models is very limited, it seems at least questionable that substantial progress will be made with this in the near future. AI research and developement thus lead us back to core philosophical questions: What is abstraction? Are there proper (Platonic) abstract objects? etc.

Somewhat surprisingly, we find ourselves on what appears to be a new path among many, leading to the realization that there is something very 'real' about abstract objects. (And that thus a lot of new problems are indeed old ones.)

Old philosophical problems related to the generation and handling of world models are for instance the 'symbol grounding problem' of the unclear connection between data and meaning,(REFs) the 'framing problem' of the unclear implementation of common sense knowledge and know-how,(REFs) the 'hard problem' of phenomenological experience, the problem of how (broad and stable) abstraction works,(REFs) and the problem of – especially also the biological implementation of – universals.(REFs) In all these cases, 'meaning' seems to evade a purely relational definition, and seems to point to the existence of qualities, forms, universals, or the like.

Quantitative and qualitative Information

A here very helpful distinction from the Philosophy of Information is the one between quantitative vs qualitative information, where the first is physically implemented and relationally defined, while the second corresponds to 'meaning' in the (original) sense of Freges 'Sinn'.(REF) And while the first corresponds to our scientific and technological understanding of information, the later is usually what we are concerned with not only in the humanities and beyond, but also in our everyday lives – and the exact connection between the two remains to be the subject of (not only) philosophical discussion.(REFs, esp. Floridi!)

However, both sub-symbolic, as well as symbolic AI run on quantitative information only, though on different levels; the problem of interfacing them is connected to the (arbitrarily) different relational assignment of 'functional' meaning to the respective (sub-)symbols. Human $brains$, too, seem to run on quantitative information and seem to work to a good extend like DNNs, with functionally understood 'meaning' as quantitatively optimized, relational structure. But at some point above(?) this, human $minds$ seem to work with bits of meaning (most basic: qualia), i.e. qualitative information.

The core problem into which we run with the idea that human minds run on qualitative information is that our current scientific world view does not allow for such bits of information to be real beyond (Humean) supervenience on material constellations. [1] If such supervenience is all

[1] It is sometimes claimed that not supervenience, but entailment is the actual problem, as 'obviously' all mental activity somehow supervenes on neural activity. Note that this is a self-fulling prophecy, as of course proper

the world allows for, it seems unclear how qualitative information could be implemented in human brains, being no more than more complicated biological DNNs. This again is the question of how complicated (now explicitly qualitative) world models could be imagined on top of DNNs.

There is of course no principal hindrance to suggest new scientific models that would allow for bits of qualitative meaning to be real beyond supervenience on material constellations, but it should be clear that this would require a major renovation not necessarily of the details of our best scientific theories, but very much so for the way we interpret their ontological implications. Such major reinterpretations are nevertheless not completely foreign to science and have indeed marked phases of most pronounced progress (think of classical and then quantum mechanics). They will nevertheless be accepted only if supported by truly impressive new experiments. This is again no principle hurdle to formulating such models, which will in any case most often aim at substantial changes not only in theory, but also in application, but as we then leave the common ground of our current understanding of science, very fundamental problems will have to be solved first.

A central problem for any such approach is to make a workable conceptualization of those bits of qualitative meaning that we want to assume to be 'properly real'. Gödel's statement 'The basis of everything is meaningful predication ... Husserl had this, Hegel did not have this: that is why his philosophy lacks clarity ...'[1] seems to point to exactly this fact. The hypothesis in the following is that the most straightforward way to conceptualize qualitative information is via distinct, non-physical properties, which then give rise to abstract entities in the same way in which physical properties give rise to physical objects. (Note that the question of the pre-existence of complex abstract objects is secondary to this; we will come back to it later on.)

Regardless of the – in the end empirical – question whether this conceptualization is indeed the best one, it seems essential to commit to a specific conceptualization. Indeed, many proposed solutions to the mind/body problem suffer from failing to do so, leading to models with some sort of diffuse mentality somehow hoovering over the material world – models which offer no way forward for drafting detailed, testable hypotheses about processes at the Neuroscience/Psychology interface.

dualists/panpsychists/idealists question exactly the assumption that this is true for all(!) mental activity.

Abstract entities

So while we usually argue (with Hume, Lewis, ...) for some sort of Humean supervenience and would understand abstract objects as regular patterns found in concrete objects, we would now (with Plato, Leibniz, Peirce, ...) understand non-physical properties – and as a result also abstract objects – as independently real. And while the former easily connects to current AI research, common theories of consciousness, connectionist neuroscience, and materialist frameworks in general (what we could overall call a 'deflated view of human thought'), the later would require (higher-level) human thinking to manipulate actual (abstract) objects, thereby suggesting that our world models are then some form of complex, brain-anchored abstract object, which in turn would require a complex sub-consciousness to 'translate' between mind and brain.

Being born in our time, understanding abstract objects as patterns seems the much more 'natural' option, but only until we come back to the question of how this would be biologically implemented and cover qualitative information. On the other hand, being allowed to understand abstract objects as bundles of universal, qualitative properties (carrying all sorts of 'meaning' that can otherwise not be fully captured, e.g. relationally via a web of words, possible worlds etc.), would be a very powerful option of any extended model. In a sense, we want both to be true; we want brains working with patterns and minds working with abstract entities. The core question really is whether this can work out scientifically.

One important aspect of this is how to differentiate and interface abstract and concrete objects in any such model. Reframed in the language of computational metaphysics, from the logical side, we can make use of (a modernized understanding of) Mally's distinction of exemplification vs encoding,(REF) and then base our investigations on object theory as worked out by Zalta $et\ al.$.[2] But from the scientific point of view, this still leaves open the question of how concrete objects (people) process abstract ones ('ideas'), i.e. the question of how to scientifically bridge this – effectively dualist – distinction.

Mally's proposal is clear from a logical point of view, but not with respect to its neuroscientific implications: If encoding is realized as physical information processing, how does this not run into the 'old' problems mentioned above? (Qualia seem to be exemplifications, not encodings, etc.) If encoding is taken to be related to proper Platonic abstract objects (a view closer to Meinong than Mally),(REF) then exemplification seems to be based on additional causal couplings that are in

need of scientific explanation.

A-world

In a recent book,[3] I have outlined a specific proposal for how to solve this $emanation\ problem$ of how to recover modern science $in\ detail$ within models that take Platonic abstract objects to exist. More generally, the book outlines which problems have to be solved and what options are available to do so. In accordance with its historical predecessors (Plato, Leibniz, Peirce ...), the resulting Model A or 'A-world' ends up as a monadic idealism, first of all to to avoid any interaction and constitution/combination problems. This way, it puts all pressure on science to fit in with this setup, acknowledging the fact that science has proven to be very adapt at re-inventing its ontology, and thereby making the whole project – scientifically – workable.

A central challenge is then to recover our current understanding of the physical micro-structure of our world, where any helpful model has to make specific ontological commitments concerning questions like what is matter? wherefrom material causation? etc. In this sense, Model A can be understood as an 'Update' of Plato and Leibniz for modern science: It assumes reality to be fundamentally build not from particles or fields in spacetime, but more generally from universal, qualitative properties. As defense, it shows how this assumption is compatible with science down to the mathematics of our best scientific theories. As offense move, the consequences are thought through to devise new, scientific experiments to evaluate the proposal. Model A is thus a scientific approach to (then pragmatic) idealism.

It can furthermore be understood as a 'consensus model', that does not take over the proposed solutions, but tries to satisfies all underlying intentions and/or intuitions of what we often consider to be mutually exclusive world-views:[4] Starting from an atomistic idealistic setup of (essentially monist) properties and a population of extremely simple 'core subjects'/monads as building blocks, i.e. an effective panpsychism, the model proposes a divergent early development towards a mind matter gap, i.e. an effective dualism, with a material world in line with current physics, i.e. an effective scientific realism, and then a gradual physical, biological and then cultural evolution of human thought.

New scientific experiments are conceivable based on such models for instance from the effects we expect from the universal nature of the building blocks, which is shown to imply a quantum na-

ture of reality at the micro-scale and relativity as a necessity for identity preservation in a world of universals. Other experiments are most likely derivable based on possible advantages for information processing with universals, and maybe for the stable organization of biological processes via non-physical feedback markers. Information processing advantages should then be coupled to specific processes and structures across the Neuroscience/Psychology interface, which should in turn show distinctive signatures in the underlying neural processes.

We would have to assume for instance that parts of the physical memory of our brain would work as no more than registers for non-physical structures, thereby implying that content richness would be somewhat independent from physical memory demand. We would also expect to see a certain level of 'super-coordination' of neural activity, and in some cases also completely missing neural activity – though a lot of constraints on this would still have to be observed for stable biological functionality.

Computing with qualia

If humans do indeed process qualitative information, then this should indeed be connected to evolutionary advantages over purely quantitative information use, and some advantages would need to have been accessible early-on; qualia would have to have an 'evolutionary purpose'.

To map out possible advantages, we would have to specify how we understand qualitative information to be different from quantitative information. The answer to this question is itself part of the project to tackle the problem of human thought. A first working hypothesis would be that qualitative information is 'transparently' given, not relationally defined; and that it 'works' via some form of material inference/implication due to an underlying 'semantic logic' or 'semantic aesthetic' of how these bits of information are related (more on this later on).

As outlined above, the most straight-forward way to conceptualize this seems to be Platonic realism; qualia are then equivalent to one-property abstract objects and whatever is special about abstract objects would also be special about qualitative information. (Abstract objects are identical if they have the same properties, they exemplify negations, can encode 'incomplete' objects, etc.)[2] Here we can build upon the above-mentioned object theory as developed by Zalta et al., although in a stronger Platonic interpretation than proposed by the authors themselves. (Accepting the problem of qualia, but failing to commit to any such specific conceptualization of qualitative

information will not allow us to propose models detailed enough to derive experimentally testable consequences.)

Based on the assumptions above, we can hypothesize that it is the ability to process qualitative information that allows humans to avoid the problems of classical AI (and combined approaches) with generating and handling complex world models: Universal properties lead to world models that automatically contain each abstract entity only once, with instances automatically connected to 'classes' and all existing (bundling) relations automatically imported. The qualitative nature of properties would result in all conceivable negations automatically implemented. And the abstract (that is; with respect to material consistency requirements incomplete) nature of objects in such world models would lead to a 'quantum object-like' behavior, where for instance 'your cat' collapses from a 'superposition' of all conceivable cats to a specific cat once I encounter it. (The abstract/concrete distinction is orthogonal to this; the status of concrete objects would depend on them being 'backed up' by a causal structure outside of the individual mind.)

Most importantly, all these features of human world models would be implemented without additional physical computation efforts for 'binding' etc. We should thus indeed be able to distinguish the two cases with respect to brain structure and activity, as in one case world model generation and handling would require less neural structure and activity.

As a first thought experiment (outlined in the above-mentioned book already), we could compare the computational effort needed for object recognition, as a first phase of 'binding', understood as the placing of properties/objects in our world model. For (naive) symbolic AI we would have to assume some type of list search, which scales proportionally to the number of properties n and the number of identifiable objects N, i.e. overall $O(n \cdot N)$. For (naive) sub-symbolic AI, object recognition scales approximately $O(N)$, because a longer object list requires more weights. For another comparison we could look at quantum computing, where a list search via the Grover algorithm scales proportionally to n and the square root of the number of amplification steps needed to reach the required accuracy, which in turn is dependent on the length of the object list, i.e. overall $O(n \cdot \sqrt{N})$. For our envisioned case of universal qualitative property handling, the object 'setup' would depend only on the number of properties, which then automatically negates all alternatives, leading to overall $O(n)$. This would be equivalent to a Grover search, but without the (here unnecessary) amplification – and would automatically import all sorts of background knowledge.

The problem with such thought experiments is that the object setup would have to be investigated in separation from the underlying quantitative processing, which would still have to initiate the generation and delivery of qualia to our (sub-)consciousness. A second problem is that is remains unclear to which specific alternative (hitherto unknown) algorithms we would want to compare our results to. Observed patterns of activity could nevertheless point into one direction or the other. To make such experiments possible, we thus need models like A-world to map out hypothetical details to test for; the search for consequences should then be formalizable in object theory.

Neuroscience

Turning to for instance the neural processing of visual information, we can check whether Neuroscience does at all allow for the existence of hypothetical activity patterns connected to the opportunity to make use of abstract entities in human thought.(REFs)

For this we find that although structures as well as pathways are known, the higher up involvement of areas for attention, memory, decision making, etc. is very complex and not well understood. Involved structures include photoreceptors, bipolar and ganglion cells, areas like LGN and then V1 to V5, etc.(REFs) Known pathways are the semi-automatic dorsal stream or 'where/how pathway', which is connected to the processing of spatial relationships, object locations, motion and in general 'vision-for-action', as well as the 'controlled' ventral stream or 'what pathway', which is connected to the processing of objects details and in general 'vision-for-perception'.(REFs) The prefrontal cortex, generally connected to working memory and decision-making, seems necessary for maintaining and manipulating objects, i.e. decision-making seems to be a critical component of object recognition.(REFs) The temporal lobes, generally connected to long term memory and semantic processing, seems necessary for contextual information.(REFs)

Neuroscience can thus pin down brain areas and pathways that are sensitive to certain object details (including positioning and motion), but can often not pin down the details (for instance concerning sub-structures or sub-pathways), and can usually not say how exactly the parts work together (especially if cross-modal).(REFs) Neuroscience can furthermore no yet explain the invariance and robustness of object recognition, the details of the influence of top-down processes, the reasons for substantial individual differences, as well as developmental details and those of in-

jury recovery.(REFs) More generally, it remains unclear how the overall distributed representations are integrated (e.g. by synchronization or convergence of neural connections), especially across modalities, and in ambiguous or conflicting cases, as well as for (with the number of features exponentially more) complex objects.(REFs)

Overall this seems to show that there is a decent amount of room for the assumptions made by model A and the hypothesis of human minds processing Platonic abstract objects. (A detailed outline of the proposed mind/brain interface can be found in the above referenced book.)

As Model A finds the micro-scale as well as higher thought dominated by the universal nature of properties, it seems tempting to look to quantum theory for new experiments in Neuroscience; similarities would be implied because objects at the bottom are not yet fully individualized, while those at the top are not anymore. (If the proposed link would hold, abstract object cognition might then serve as a blue print for the development of quantum computing algorithms, though unlike in abstract object manipulation, quantum computing could make use of only a few and sharp (material) semantic implications.)

Such experiments would include: 1. 'Neuro Bell experiments', in which brain processing statistics would indicate 'entanglement', i.e. simultaneous, distant activity related to the same content. 2. 'Neuro Kochen-Specker experiments', in which brain processing statistics would indicate 'contextuality', i.e. identical activities delivering different contents. 3. 'Neuro double slit experiments', in which brain processing statistics would indicate nonlocal interference patterns without experiments being able to pin down the 'trajectories' of neural activity. Many more experiments might be in store here. It is important to note here that in Model A such effects would not be expected because of the assertion of long-range quantum effects in the brain, or because consciousness would be somehow intimately connected to quantum effects – neither option is considered in Model A. Insted, the core feature would always be that besides physical also non-physical elements would be available for human thought, so that when we look at the brain, we would observe only parts of the overall activity.

More specific suggestions for experiments require a more detailed description of the proposed Neuroscience/Psychology interface; I hope to be able to deliver first suggestions for this in the not too distant future.

Questions & Answers:

Does the ontology of such models include an infinity of basic properties? First of all, in Model A, this is an empirical question, much like how we do not know what physical properties we should take for granted before experimental investigations. In addition, seemingly infinite 'ranges' of properties (color shades etc.) might already be the result of the bundling of underlying properties (color and shading etc.). But most importantly; Platonic realists (and many scientists implicitly) are usually indeed willing to accept an infinity of at least basic mathematical entities already.

Do all abstract objects pre-exist in such models? The starting assumption in Model A would be that non-physical, basic properties, but not complex abstract entities are pre-existing. Complex entities are formed in individualized minds based on physical signals being converted to non-physical properties and then simple abstract objects by our sub-consciousness. Against the backdrop of our existing world model, physical signals would have to be helpful enough to enable our sub-consciousness to generate fitting simple abstract objects.

At the lower level, our shared biology would guarantee (at least partially) successful communication, at the higher, proper abstract level, the universality of the basic building blocks would have to lead to the seemingly universal nature of also complex abstract objects.

Additional arguments available, one could still argue for the pre-existence of all abstract objects; this would be connected to the question of whether rational optimism holds; see next paragraph.

How does the proposed pragmatic idealism differ from rational optimism, for instance in the footsteps of Gödel? Both pragmatic idealism and rational optimism assume conceptual and mathematical realism and take physical (unlike subjective) time to be not fundamental, but the outcome of the development of a material part of reality, i.e. of a subset of objects with properties which change according to consistency rules. Both allow for 'ontological arguments', but in pragmatic idealism, these are only 'trivially' correct and without immediate necessary consequences for the material world, as they prove abstract existence only.

The core difference is that for rational optimism, the pre-supposed (wished-for?) effectiveness of logic to deal with ontology requires necessary implications ('strict relations') between properties, i.e. 'epistemic stability' as L. Hornischer calls it,[5] which leads to all sorts of 'eternal truths'; Hegel's semantic logic made proper, so to say. We can then identify positive properties, that existence is positive, etc. This way, all contingency, including its yardstick physical time, ends up to be just down to our limited experience of a perfectly well-ordered whole, in turn implying theism.

In pragmatic idealism, the idealistic nature of reality is not a rationally derivable outcome, but an assumed fact that scientific inquiry might establish as best overall fit. This allows for a contingent understanding of positive properties, existence, etc. Ontological arguments then seem to imply at the meta level that if there is a 'semantic logic' (instead of a 'semantic aesthetic') between properties, then all abstract entities are (pre-)existing – which would take us back all the way to Plato's Meno.

But pragmatic idealism allows for the implications between properties to be of 'aesthetic' nature, i.e. much broader, although that might be the case only beyond certain property classes like those of mathematics. Rational optimism would then be justified only if all implications are actually narrow, not broad, which does not seem to be the case already for colors.

How would imagination and perception differ? First of all again, in Model A, this is an empirical question, much like how we do not know physical facts before experimental investigations. Secondly, the boundary between imagination and perception does not seem to be as rigid as we often think; phenomena like substance abuse, schizophrenia, or the imagination of 'proper physical' objects in meditation show this. Any useful model should nevertheless have a hypothesis for the important observation that normally imagination and perception differ.

In Model A, a straightforward hypothesis would be that the two happen at different levels of abstraction. While in perception qualia properties would be pushed by our sub-consciousness into our consciousness, imagination would most often happen at a level where more easily manageable concepts stand in for the complex objects of perception. The necessity of such a distinction between qualia and concepts of qualia was anyhow proposed in Philosophy by Nida-Rümelin.?? 'Proper' encoding and abstraction in the common sense would be found at the higher level(s). The activation of perception in imagination would nevertheless remain a possibilitiy, e.g. in those

special cases mentioned above.

The even more interesting question would then be, why and how such a layered mental structure would have evolved. A first hypothesis would be the following: The direct handling of perceptual objects would always take them and their implications as true; this would work like what we expect for instincts. (Zoos are for instance most often not a fun trip for animals, as certain reactions implied by certain perceptual objects cannot be avoided by them.) Being able to push qualitative signals to a more abstract level would take pressure from lower levels to immediately follow implications, which would later on allow for reflection and (attempts at) global optimization; once such a mechanism is available, higher-level abstraction is in reach.

One powerful feature of language would then be to indicate 'here is a thing', without attaching any (or more than the minimum?) of properties to it. (Large parts of?) abstract thought would then indeed be conceptual, but taking all of human thought as language-based would still not be right.

The hypothesis above might be experimentally testable: Do ('lower') animals remember individuals? Do we not remember individuals in perceptual detail? Is the ability to control impulsivity somewhat correlated with better abstraction skills? Etc.

Are qualia not privately generated and exclusive? In model A qualia would be treated on the same footing as physical properties: Individual material objects would form from 'particle bundles' of universal material properties; think of universal spin ending up in a concrete particle. Individual non-material objects would form from bundles of universal non-material properties, think of universal red ending up in a specific mental object.

Universal material properties are essential to our understanding of laws of nature. (Universal non-material properties might then be the basis for successful intersubjective communication beyond simple 'success in a 3D material world' cases.) It is very hard to conceive science-like models of phenomenological experience without universal non-material properties: This would correspond to a view of physics in which each particle is having its own spin-property, so that we would have to check for every particle pair according to which pair-specific law their individual spin-properties would interact.

We are happy to embrace Platonism for physical properties (universal mass, spin, etc. properties for all particles) – the idea here is that Platonism should be our first best guess for mental

properties, too.

How would our minds be anchored to ensure their identity in a world of universals? Models like 'A-world' are essentially extrapolating down from biology, instead of up from physics; we know of subjects and qualities at the (human) meso-scale, have good reasons to believe that things at the (e.g. animal) sub-meso-scale look similar and propose that this is the case all the way down to elementary particle physics.

Starting from universal properties and the evolution of a monad population only, we then observe an easy loss of identity at the micro-scale, because any conceivable particles are abstract objects as long as their properties are not handled consistently. Spatial properties are one way out; functional space would develop to ensure identity, not the other way around. Stable identities can then form upwards from the molecular scale thanks to evolutionary learned consistency rules for the handling of the subset of properties that make up the physical world .

To maintain their stable identity, when monads reach out to non-physical properties, these must be 'anchored' in individualized bodies for stable growth (scepticism and truth become issues then). Only complex minds do then allow for the 'free' handling of abstract objects (anchored in a world model, in turn anchored in a body, etc.). And such minds would have to have a complex (and multi-monadic) sub-consciousness to handle both physical and non-physical properties by different rules. This would lead to additional 'neural design principles' (like energy optimization), for instance the requirement to 'pass everything vitally important through matter'. (As only then the 'stability' of biological processes is guaranteed by the consistency of how material properties are handled.)

More details on such evolutionary monad models can be found in the book referenced above. Alternatives, which nevertheless most people will consider at least partly outdated, are for instance Leibniz' pre-established harmony or Plato's dispositional/causal properties. The presumed existence of Platonic abstract objects seems to require some form of monadic idealism; the only class of models not only compatible with modern science, but at least potentially helpful for the explanation of phenomena at the Neuroscience/Psychology, seem to be the one of evolutionary models. A scientific formulation of such models seems to possible; whether this is of any relevance will have to be judged based on whether experimental evidence for or against such models can be found.

References

[1] H. Wang, A Logical Journey: From Gödel to Philosophy, MIT press 1996, p 168.

[2] E.N. Zalta, D. Kirchner, U. Nodelman, Principia Logico-Metaphysica (Draft/Excerpt) as of May 16, 2025. https://mally.stanford.edu/principia.pdf

[3] M. Korth, Information, Intelligenz und Idealismus, Brill/Mentis 2026 (upcoming); earlier Englisch manuscript 2025. https://philpapers.org/rec/KORIIA-4

[4] M. Korth, How to construct consensus models to (maybe) make sense of the mind-body problem, 2025. https://philpapers.org/rec/KORHTC

[5] L. Hornischer, A Stability Interpretation of Gödel, contribution to the 2025 Kurt Gödel Essay Prize. https://kurtgoedel.de/cms-83FO/wp-content/uploads/2025/07/M-Stability-Interpretation-of-Goedel.pdf