Quantification, Inference, and Ontology

GABRIEL UZQUIANO

Thomas Hofweber has written a very rich book. In line with the conviction that ontology should be informed by linguistic considerations, he develops a systematic approach to central ontological questions as they arise in different regions of discourse. More generally, the book seeks to cast light upon the nature of ontology and its proper place in inquiry. His preferred methodology is not without consequence: it promises, for example, to solve what otherwise look like intractable philosophical puzzles raised by arithmetical practice and numerical discourse, and, likewise, his treatment of propositional discourse has ramifications for larger questions to do with the prospects of metaphysical inquiry generally.

   In what follows, I would like to comment on the general approach to ontology and to look at the special case of propositional discourse and its ramifications for metaphysics.

1. Quantification and Ontology

At the heart of the book is the thesis that quantifiers are semantically underspecified. There is at least a distinction between an external and an internal reading of a quantifier. The external reading is tied to the representational role of the quantifier whereby the world is portrayed to come with a domain of objects over which the quantifier ranges. On the external reading, the truth conditions of a quantified sentence are given in terms of a condition on an external domain of individual objects: a sentence of the form 'something is *A*' is true if, and only if, some member of the external domain is in the extension of A. Other quantifiers are associated with different conditions on the external domain.

      On the internal reading, however, a quantifier is governed by its inferential role. The quantifier 'something', for example, is governed by a rule of existential generalization, which allows it to act much like a trace for its substitution instances. By a substitution instance, however, Hofweber doesn't merely mean a referential expression – a determiner like 'two' is, by his lights, an important counterexample to the identification – nor merely an expression that can occupy the same syntactic position as the quantifier – a quantifier phrase like 'nothing' would likewise refute this criterion.[1]

Of course to make the inferential role of the quantifier precise, we should specify the class of appropriate substitution instances. This is quite a substantive task, but we can in the meantime, make do with our ordinary judgments as to what constitutes a valid instance of existential generalization. What's important for now is the observation that on the internal reading of the quantifier, the specification of the truth conditions of a sentence like 'something is A' require no mention of an external domain. Instead, they boil down to the truth conditions of a sentence of what would be an infinitary extension of a fragment of English: '$\bigvee_{t \in \Sigma} t \ is \ A$', where Σ is the set of appropriate substitutions for the quantifier. The account is different for other quantifiers: the truth conditions of a universally quantified sentence 'everything is A' come down to the truth conditions of an infinite conjunction of the form '$\bigwedge_{t \in \Sigma} t \ is \ A$'. Much more delicate is the case of generalized quantifiers like 'most A are B' or 'few A are B', but they too are covered by the proposal.[2]

The focus on the unary existential quantifier is far from coincidental. The primary ontological question for numerical discourse, for example, is whether something is a number on the external reading of the existential quantifier. And similarly for many other basic ontological questions in other areas of inquiry. One would like to know how best to tackle these questions, and more generally, whether they fall within the remit of metaphysics or are to be settled by other parts of inquiry.

The distinction between the internal and external reading of the quantifiers informs Hofweber's methodology. He recommends we start with the question of whether the quantifiers are generally read internally or externally in a given area of discourse. The answer to this question determines an answer to the primary ontological question for the relevant domain as well as a verdict as to whether it falls within the remit of metaphysics or some other part of inquiry. Numerical discourse is a case in point. Hofweber argues that as generally used in arithmetic and ordinary discourse, numerical terms are not referential expressions and numerical quantifiers are read internally and not externally: they don't purport to range over an external domain of objects. Moreover, internalism for numerical discourse entails a negative answer to the primary ontological question of whether something is a number – on the external reading of the quantifier. This question falls within the remit of metaphysics but it's guaranteed to have a negative answer.

The internalist program for numerical discourse yields a solution to what Hofweber calls *the central tension* in the philosophy of arithmetic. That's the problem of reconciling the objectivity of arithmetic and the fact that it seems to enjoy a subject matter of its own with the fact that thought alone appears to suffice to make progress in arithmetic. Internalism provides a solution to this problem: one grants that natural numbers provide a subject matter for arithmetic and nevertheless disavows any reading of this claim on which natural numbers are themselves objects in the external domain.

But versions of the central tension arise for other areas of mathematical discourse, and there is no guarantee that they can – or should – be settled by similar means. The arguments for numerical

internalism are quite subtle and depend on very specific linguistic evidence that may not be available in other areas such as real analysis and set theory. For all we know, externalism could be true of talk of real numbers and sets in which case we would need a separate approach to the relevant variation on the central tension in each case. Whatever the case, Hofweber rejects the presumption that we should expect a uniform solution to all these puzzles.

Internalism, after all, is not a blanket approach to ontological questions. Hofweber (2016, Chapter 7) is careful to develop the case for the external reading of quantification over material objects in ordinary discourse. This means, for him, that the basic ontological question of whether something is a material object is not one we can settle from the armchair; it is an empirical question to be settled by other parts of inquiry.

In sum, Hofweber advocates a *language first* approach to ontology: if you want to solve the primary ontological question for a given domain of discourse, figure out, first, whether internalism or externalism is true of that domain. Careful reflection on the available linguistic evidence should guide the question of whether any appropriate substitutions for the quantifiers are indeed referential expressions. If the answer is 'no' and quantifiers are generally read internally, then the answer to the primary ontological question falls under the remit of metaphysics and is guaranteed to be 'no'; otherwise, if the quantifiers are generally read externally, then the answer to the primary ontological question is no longer guaranteed to be 'no' but it will fall under the remit of other areas of inquiry.

The primacy of linguistic considerations and the fact that central ontological questions are partly individuated by the language in which they are couched place important limits on the scope of the answers they receive. Even if one answers the primary ontological question of whether there are numbers – as phrased in English, there are, by Hofweber's lights, significant limits to what one, may have accomplished at this point. For he is careful to explain that a negative answer to the ontological question as to whether something is a number provides little or no support for nominalism. Forget whether numerical expressions are referential as used in arithmetic and ordinary numerical discourse, maybe there are von Neumann ordinals or abstract positions in an ω-structure; these abstract objects would be *numberlike* in that they would enjoy many of the central features externalists associate with natural numbers. And there may even be a language other than English that comes equipped with words that are used referentially to pick them out. Even if externalism turned out to be true of talk of these *numberlike* objects in the new hypothetical language, this observation would be quite irrelevant for the question of whether something is a number as phrased in English. Whether the external domain of objects contains *numberlike* objects with some of the central features externalist associate with numbers is beyond the scope of the primary ontological question associated with arithmetical discourse.

In what remains, I would like to focus on the case of propositional discourse as another case study, one with ramifications for even larger metaphysical questions.

2. Internalism and Propositional Discourse

The book touches on the ontological question of whether there are propositions and its ramifications for the ambitions of metaphysics. On the approach outlined in the book, the crucial issue is whether internalism or externalism is true of propositional discourse. If externalism is true, then that-clauses refer to propositions to be found in the external domain of objects. Internalism, on the other hand, insists on an internal reading of quantification over propositions.

Some philosophers take externalism to provide the best explanation of the validity of simple quantificational inferences. Consider, for example:

(1) Alice believes that snow is white.

(2) Alice believes something.

They construe a that-clause like 'that snow is white' as a referential term for a proposition; they take attitude verbs like 'believe' to involve a binary relation between an agent and a proposition; and they conceive of the inference from (1) to (2) as an instance of existential generalization whose conclusion requires the external domain to contain at least one proposition. They presuppose, in other words, an *external* reading of the quantifier.

Hofweber has no quarrel with the validity of the inference but he objects to the externalist gloss of the conclusion. On the *internal* reading of the quantifier, the truth of the conclusion places no demands on reality over and above the truth of a sentence of the form 'Alice believes that p' where 'p' is replaced for a sentence of the language. More generally, the truth conditions of 'Alice believes something' are equivalent to the truth conditions of an infinitary disjunction of the form '$\bigvee_p Alice\ believes\ [that\ p]$' for each replacement of p for a sentence of our language. Since (1) is one of these disjuncts, we have that (2) logically follows from (1). In other words, the internal reading of the quantifier provides a rival account of the validity of the quantificational inferences in question.

Hofweber (2016, Chapter 8) breaks the stand off by noting that on balance, reflection on the available linguistic evidence suggests that that-clauses are in fact not referential expressions, and whatever else their semantic role may be, it is not to refer to a proposition. So, Hofweber suggests, internalism is the only game in town.

This makes the central ontological question of whether something is a proposition – read externally – to be part and parcel of metaphysics, but at the same time according to Hofweber, it guarantees a negative answer to it: there are no propositions to be found in the external domain of objects. This is not to deny that there are hard metaphysical questions in the vicinity. Internalism is not meant to help us, for example, with questions of fineness of grain: are necessarily equivalent propositions one and the same? And if facts are true propositions, we still face the question of how fine

grained reality itself is. These, and similar questions, require a separate treatment and remain as difficult for the internalist as they are for the externalist.

There's a hitch. Internalism appears to face a formidable objection.[3] No matter how expressively rich one may take English to be, few might be inclined to think that it is rich enough to express each and every proposition. Some propositions, you may think, are not expressible in English much like it seems to us, for example, that many propositions we routinely express in English could have not been expressed by a sentence of ancient Greek. And yet, (3) below comes out false when the quantifier is read internally:

(3) Something is not expressible in English.

Internalists read (3) to be equivalent to an infinitary disjunction:

(4) $\bigvee_p [that\ p]\ is\ not\ expressible\ in\ English,$

whose instances result from a substitution of an English sentence for 'p'. Since no such instance can be true, the disjunction comes out false. It would seem a mistake to conclude that English is expressively complete in ways in which more impoverished languages may not be. For a parallel argument would have been available to a speaker of ancient Greek, even though, in hindsight, it may seem to us that they lacked the resources to express the proposition that Alpha Centauri is a constellation or that there are quarks. What gives?

Many sentences of English contain context-sensitive expressions, which make a difference to what an utterance of the sentence by a given speaker may express relative to one context or another. This suggests a distinction between what can be expressed by an utterance of a context-independent sentence in a given language and what can be expressed by an utterance of a context-sensitive sentence of the language at a context. Even the latter calls for a distinction between what can be expressed by a context-sensitive sentence of the language at an actual context and what can be expressed at one context or another – regardless whether any speakers of the language ever find themselves in them. This difference motivates a distinction between three conceptions of expressibility, which correspond to at least three different theses in the vicinity of (4). The first two aren't very plausible at all for a language like ancient Greek – or even English: there's no reason to assume ancient Greek contains a context-independent sentence a speaker could use to express the proposition that Alpha Centauri is a constellation.[4] Nor is it clear ancient Greek speakers ever found themselves in a position to use a context-sensitive sentence to express the proposition that Alpha Centauri is a constellation. No matter, Hofweber thinks, there *could* have still been a context in which Alpha Centauri is referred to by a context-sensitive expression by a speaker of ancient Greek.

In order to accommodate quantification over propositions – whether merely expressible by utterances of context-independent or expressible by utterances of context-sensitive sentences relative to one context or another – Hofweber models the truth conditions of quantified sentences in an

interpreted formal language equipped with external quantification. The truth conditions of an existentially quantified sentence of the form (5) are given by (6) – where the outer quantifiers are read externally to range over eligible referents for context-sensitive expressions of the target language:

(5) Some proposition is A.

(6) $\exists \vec{x} \, \bigvee_p [that \; p(\vec{x})]$ is A.

In particular, (3) becomes:

(7) $\exists \vec{x} \, \bigvee_p [that \; p(\vec{x})]$ is not expressible in English.

The claim is true on the first two conceptions of expressibility and false on the third one. But whatever confidence we have in the falsity of (3) should no longer be undermined by the apparent inability to express the proposition that Alpha Centauri is a constellation by a context-independent sentence in a language like ancient Greek – or the fact that no ancient Greek speakers may have found themselves in a position to use a context-sensitive sentence to express that proposition. The story is more complicated for the proposition that there are quarks, but there is no reason to think it cannot be given.

Notice, however, that we need a more discriminating conception of expressibility for (3) to draw a real contrast between internalism and externalism. Externalists, after all, take apparent quantification over propositions to be restricted quantification over the external domain. Since propositions are objects, they too are eligible referents for context-sensitive expressions at a context, and whether or not an English speaker ever finds herself in a context at which a given proposition is referred to by a context-sensitive expression, there's no reason to think there couldn't be such a context. One may perhaps question whether the mere fact that a proposition p is eligible to be the referent of a context-sensitive expression of the language suffices for the proposition to be expressible. For one may again require us to find a sentence of the language to express p at a context. But unless propositions are very fine grained indeed, it would seem that an utterance of the sentence 'that is the case' at a carefully chosen context should be able to express one and the same as proposition as p.[5]

Internalism helps with the question of whether we can, in fact, expect our metaphysical theorizing to proceed unconstrained from any representational limitations on our part. Facts are correlated to true propositions – indeed, some take facts to be true propositions. If you take reality to consist of facts, you may naturally expect some of them to lie beyond our ken. But as Hofweber is careful to note, we must fix what counts as a relevant representational limitation in order to properly asses this hypothesis. And he opts for what he calls an *object permitting* conception of ineffability: whether a proposition – or fact – is ineffable cannot be a matter of our inability to represent an object in the external domain. On this understanding of ineffability, it follows from internalism that there are no limits to what we can represent in thought and language.

Hofweber (2016) uses this observation to allay *the dark vision* according to which we lack representational access to large-scale features of reality that lie beyond our ken. Since there are no

limits to what we can represent in thought and language, there are no restrictions on the ambitions of metaphysics. This is again supposed to mark a real contrast between internalism and externalism: externalists take that-clauses to be referential and quantification into their position to range over propositions conceived as objects in the external domain. If the external domain is, as is plausible to assume, independent of thought and language, externalists shouldn't expect to be able to represent every proposition in thought and language. There's no reason to think that a finite mind like ours, which developed in a certain environment as a result of certain selection pressures, should be able to represent any fact whatever. Instead, the default assumption should be the ineffability thesis according to which some propositional aspects of reality are, in fact, ineffable.

Externalists, according to Hofweber, should approach metaphysics with a dose of modesty and humility. They should think there are severe limits on what metaphysics can accomplish. But this again requires a more discriminating conception of ineffability. On the object-permitting conception of ineffability, there mere fact that we cannot represent an object in the external domain is never a relevant limitation. Since propositions – and facts – are, for the externalist, objects in the external domain, it looks like we can, for purposes of assessing the ineffability thesis, grant that we can have a representation of them. Now, to be sure, the question of whether a fact is ineffable is not the question of whether it is ever an eligible referent for an expression at a context but rather whether we can utter a sentence that expresses, states, or represents that fact. But whatever the nature of the fact, if we can refer to it, by the lights of externalists, we can presumably represent the proposition that *it* is the case.

3.   Propositional Externalism Revisited

Whatever you make of the internalist program, the problem, according to Hofweber, is that there is no real alternative to it when it comes to propositional discourse – or at least not if that-clauses are not referential terms. Let us grant that on balance, the linguistic evidence provides sufficient support for this claim. [6] We now explore the question of whether there is an alternative construal of the external reading of the quantifier in (2), one on which it would still have a representational role, though not one that requires it to range over any *objects* in an external domain.

What I have in mind is the type-theoretic perspective inspired by the work of Frege and Church and, more recently, advocated by philosophers like Tim Williamson. On their picture, there is an alignment between certain grammatical categories and the type-theoretic hierarchy: singular terms are expressions of type e. An expression of type $(t_1, \ldots, t_n)$ combines with expressions of types t_1, ..., t_n, respectively, to form a sentence. A sentence is itself an expression of type $()$, whereas some monadic predicates combining with expressions of type e to form sentences is an expression of type (e). Quantifiers come in different types: if a unary quantifier combines with an expression of type (e) to form

a sentence – much like in 'something is red' – it has type $((e))$. But if the unary quantifier combines with an expression of type $(())$ to form a sentence – much like in 'something is true' – it receives type $((()))$.

On a set-theoretic semantics for the theory of types, expressions of type $(t_1, \ldots, t_n)$ are assigned *sets* of n-tuples of semantic values for expressions of types $t_1, \ldots, t_n$, respectively. This is not, however, the interpretation these philosophers have in mind; instead, they take expressions of type $(t_1, \ldots, t_n)$ to correspond to irreducible n-ary relations whose arguments correspond to expressions of types $t_1, \ldots, t_n$, respectively. On this picture, propositions are 0-ary relations that correspond to sentences of type $()$, predicates of type $(())$ combine with them to form sentences, and unary propositional quantifiers are expressions of type $((()))$.

These philosophers are of course careful to distinguish quantification into different grammatical positions, and they deny that quantification into predicate or sentence position – read externally – is quantification over a domain of objects. Predicates, after all, do not refer, they predicate, and, the thought goes, it would be a mistake to understand quantification into predicate position as quantification over a domain of objects to which predicates refer. Likewise, for quantification into sentence position. This tradition grants a representational role to all these quantifiers in that they take the world to be portrayed to come with an external domain in the form of a type-theoretic hierarchy of objects and irreducible relations of different types.

There's a familiar expressive problem for the position: one may be tempted to describe these relations as the semantic values of expressions of different categories. But now one may object: what could these semantic values be, if not objects? A better characterization of the type-theoretic hierarchy would need to appeal to the very expressive resources we want to interpret.[7] But whatever force the objection may have, what matters for present purposes is that the position opens the way for an externalist reading of propositional quantification on which that-clauses needn't be referential.

Hofweber (2016) briefly mentions similar moves in the vicinity but objects that it's not enough to say that these quantifiers are primitive and that they cannot be spelled out other than in terms of themselves. If there is a difference between all these quantifiers, one should explain what it is and how it emerges in the first place.

That's a tall order of course. But one could perhaps begin with the observation is that there is an inferential characterization of the external reading of the existential quantifier of type (e) on which it corresponds to a restriction on the internal reading. This may seem in tension with the claim that a sentence of the form 'something is A', read externally, can be true even if no sentence of the form 't is A' is true for any substitution for the quantifier.[8] However, the distinction between context-independent and context-dependent instances of a quantifier can be used to navigate this obstacle. The key again is to make a distinction between what it is for a context-independent expression to refer to an object in an external domain and what it is for a speaker to use a context-sensitive expression to refer to it relative to

a context. There is, moreover, a distinction between what a context-sensitive expression refers to in actual contexts, and what it could refer to in a context – regardless whether any actual speakers of the language ever find themselves in it.

The external reading of the existential quantifier of type $((e))$ is governed by the rule of existential generalization, which is now implicitly relativized to a context – whether actual or not:

- *from 't is A', infer 'something is A',*

where t is a referential term of type e, which successfully performs its referential function relative to the context. And indeed, the external reading of the unary existential quantifier appears to align with the corresponding restriction of the internal reading of the quantifier.

What's a *referential term*? It's a linguistic expression whose primary semantic function is to refer to an object. Names are paradigmatic referential terms, but other expressions may play the role as well. Moreover, some of them may play the referential role only relative to a context of utterance. A demonstrative, for example, can refer to an object relative to a context of utterance. The appeal to context is supposed to help with the case in which some object is indeed A even if nothing that is A is ever denoted by any context-insensitive referential term of the language. For if something is indeed A, then there could still be a context and a context-sensitive expression that would denote an object that is A relative to it. On this approach, the external reading of the existential quantifier of type $((e))$ corresponds to a proper restriction of its internal reading, one which restricts the range of eligible substitutions to referential terms – some of which may be context-sensitive – which successfully performs their referential function relative to a context, whether actual or not.

The hope now would be to motivate the external reading of quantification into other grammatical positions as aligned with a certain restriction of the internal reading, one on which we restrict the range of eligible substitutions for the quantifier to expressions of the relevant grammatical category that successfully perform their primary semantic function relative to a context, regardless of whether any actual speakers ever find themselves in it.[9] Sentences are not referential expressions, but when competent speakers encounters one, they expect it to come with a semantic value conceived here as an irreducible 0-place relation. [10] Some 0-place relations may of course be expressed by context-insensitive sentences, but a great many of them may not be. When we use a quantifier of type $((()))$ to externally quantify into sentence position, we restrict the range of eligible substitutions for the internal reading to context-sensitive and context-insensitive expressions of type $()$ that successfully perform their semantic function relative to a context, whether actual or not. On this picture, if something is A – where, very roughly, 'something' is understood as a quantifier of type $((()))$ read externally – there *could* be a context and a context-sensitive expression of type $()$, which, relative to that context, expresses a 0-place relation that is indeed A. The external reading of quantification into sentence position becomes

a device for generalization of successful semantic role for expressions of type () relative to a context, whether actual or not. And this in turn amounts to a generalization over irreducible 0-place relations.[11]

The type-theoretic approach comes with an alternative account for the validity of simple quantificational inferences. The argument from (1) to (2) above is an instance of existential generalization for the propositional quantifier. This is still compatible with different proposals for regimentation. You may take the argument to involve a sentential operator, 'Alice believes that', in which case that clauses appear to play no role in the argument, or you may nevertheless think of a that clause as a device for nominalization, which is itself not a referential expression.[12] What is important for us is that there is no need to conceive of that-clauses as referential expressions.

Unlike other formulations of propositional externalism, this version makes room for a distinction between the view of reality as a totality of objects and the view of reality as a totality of facts. Neither facts nor propositions more generally are objects. Nor there is any reason to think that the object-permitting conception of ineffability can help us come closer to being able to express that a given fact is the case. Whether this leaves us vulnerable to *the dark vision* will now depend on what conception of ineffability we use to assess the ineffability thesis.[13] But that's a topic for another occasion. The goal here has merely been to explore the prospects of an alternative construal of propositional externalism; whether the approach places any constraints on the ambitions of metaphysics is an important, though separate question.[14]

References

Hofweber, T. 2006. Inexpressible Properties and Propositions. *Oxford Studies in Metaphysics* 2, pp. 155-206.

Hofweber, T. 2016. *Ontology and the Ambitions of Metaphysics*. Oxford: Oxford University Press.

Prior, A. N. 1971. *Objects of Thought.* Oxford: Clarendon Press.

Recanati, F. 2004. 'That'-Clauses as Existential Quantifiers. *Analysis* 64: 283, 229-35.

Williamson, T. 2003. Everything. *Philosophical Perspectives* 17:1, 415-65.

Williamson, T. 2013. *Modal Logic as Metaphysics.* Oxford: Oxford University Press.

Wright, C. 2007. On Quantifying into Predicate Position: Steps Towards a New(tralist) Perspective. In *Mathematical Knowledge*, ed. M. Leng, A. Paseau, and M. Potter, Oxford: Oxford University Press, 150-74.

---

[1] Otherwise, we could move from 'nothing is non-self-identical' to 'something is non-self-identical'. The decision is harder for other expressions. Should, for example, definite descriptions like 'the author of *Ontology and the Ambitions of Metaphysics*' be eligible substitutions for *t*?

[2] Hofweber (2016) undertakes this task in appendix 3.7 in p. 168.

³ See Hofweber 2006 for further discussion of this problem.

⁴ There is nothing special about ancient Greek here. There is no reason to think that English should contain for each real number, a context-independent sentence a speaker could use to express the proposition that it is irrational at any context.

⁵ Externalists might of course object to the analysis of the truth conditions of (3) in terms of (7). The point is merely that whatever proposition in the external domain is assigned to the free variable *x* in the open sentence '*x* is not expressible in English' should itself be an eligible referent for a context-sensitive expression in English, e.g., a demonstrative like 'that'. And, moreover, it should be expressed by an utterance of a sentence of the form 'that is the case' at a carefully crafted context.

⁶ This is amply supported and documented in chapter 8 of Hofweber 2016.

⁷ This familiar expressive limitation is of course related to Frege's infamous concept horse problem. This and related concerns are discussed, for example, in chapter 5 of Williamson 2013.

⁸ This will happen whenever the truth of the external reading of the quantifier comes down to the existence of an object in the external domain of quantification that is *A* despite the fact that we lack a term for it.

⁹ Wright (2007) has argued, more generally, for a picture of external quantification as a device for generalization of semantic role.

¹⁰ For a broad outline of the idea that our grasp of sentential quantifiers is rooted in our grasp of sentences, see Williamson 2003 and, more recently, chapter 5 of Williamson 2013.

¹¹ These remarks are not meant to suggest a reductive account of the external reading of quantification into sentence position merely in terms of the internal reading. For they remain silent as to what it is for something to count as an eligible substitution of the quantifier relative to a context, and what it takes for a sentence to successfully perform its semantic function. Someone could, for example, deny that any eligible substitutions could themselves include quantification into sentence position. So, she would question that sentence like 'something is true' could ever be an eligible substitution for the internal quantifier.

¹² The first alternative is advocated by Prior 1971 and a version of the second is discussed, for example, by Recanati (2004).

¹³ Though far-fetched, I can imagine a theorist making the case for a relation-permitting conception of ineffability according to which whether a proposition is ineffable cannot merely be a matter of our inability to represent a relation in the external domain.

¹⁴ For helpful comments and discussion, thanks to Matti Eklund, Jeremy Goodman, and Thomas Hofweber.