Leszek Kołodziejczyk
Institute of Mathematics
University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
lak@mimuw.edu.pl

Warsaw, October 5, 2023

Report on Erfan Khaniki's PhD thesis

Erfan Khaniki's thesis, titled *Impossibilty* [sic] *results in Proof Complexity and Arithmetic*, consists of four essentially independent papers prefaced by a brief introduction. Three of the papers have already been published. In this report, I describe the contents of the papers, attempt an evaluation of the results, make some comments about the presentation, and state my conclusion.

A list of detailed comments on both mathematical content and presentation is attached as an appendix to this report. The comments concern all the papers, including the published ones.

Contents. The thesis is based on the following papers, all of which are single-authored:

A. *On Proof Complexity of Resolution over Polynomial Calculus*, ACM Transactions on Computational Logic, 2022.

B. *Nisan-Wigderson generators in Proof Complexity: New lower bounds*, Computational Complexity Conference, 2022.

C. *Jump operators, Interactive Proofs and Proof Complexity Generators*, preprint.

D. *Not all Kripke models of HA are locally PA*, Advances in Mathematics, 2022.

Papers A-C are about broadly understood proof complexity, with connections to nonstandard models of arithmetic and to bounded arithmetic. Paper D is about models of Heyting Arithmetic, the intuitionistic counterpart of Peano Arithmetic.

Paper A concerns the proof complexity of systems of the form $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$, that is refutation systems in which proof lines are disjunctions of bounded-degree polynomial equations over a fixed field $\mathbb{F}$ (or in some results, over a general ring rather than a field). The main theorem of the paper is a size-width connection for such systems, which says the following: given a formula F, a lower bound on the minimal width of a $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ refutation of F (the width of a refutation π being the maximal number of disjuncts in a line of π) implies a lower bound on the minimal length (number of lines) of a treelike $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ refutation of F. Furthermore, if the width lower bound is strong enough, it will also imply a nontrivial length lower bound on arbitrary (daglike) refutations of F. Such a width-size connection has been known to hold for Resolution refutations for more than 20 years, and it has been a fundamental tool in proving Resolution size lower bounds. In fact, the author proves the connection for $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ by reducing it to the one for Resolution.

As a consequence of the size-width connection for $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$, combined with a routine translation from $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ into Polynomial Calculus and with known PC degree lower bounds, the author obtains a number of exponential lower bounds for treelike $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ and the first nontrivial lower bounds on the length of daglike $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ refutations of CNF formulas (which are also the first length lower bounds of any kind on daglike $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ over finite $\mathbb{F}$). Unfortunately, the bounds for the daglike systems are not even quadratic – though it should be pointed out that already for e.g. $\mathsf{Res}(\mathsf{PC}_{1,\mathbb{F}_2})$ proving superpolynomial lower bounds is a well-known and seemingly quite difficult open problem. The technical reason why the daglike lower bounds are so modest is that the expression specifying the size-width relationship for daglike Resolution contains a term for the number of variables in the formula being refuted. The author's reduction from $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ to Resolution involves adding new auxiliary variables, so this term blows up and hence the obtained connection between $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ width and size is too weak to imply quadratic or better size lower bounds.

In Paper B, the author contributes to the study of proof complexity generators, that is $\mathsf{NP} \cap \mathsf{coNP}$ maps $g\colon \{0,1\}^n \to \{0,1\}^{m(n)}$, with $m > n$, such that for some given proof system P, proving that any specific b is in $\{0,1\}^m \setminus \mathrm{rg}(f)$ is hard for P. It has been conjectured by A. Razborov that for a sufficiently computationally hard $\mathsf{NP} \cap \mathsf{coNP}$ property f, what is known as the Nisan-Wigderson generator based on f is hard even for quite strong proof systems. The main result of the paper is that if f is symmetric (i.e. given a fixed input length, depends only on the number of 1's in the input string) but cannot be expressed by subexponential-size DNF's or CNF's, then a particular formalization of the Nisan-Wigderson generator based on f is a hard proof complexity generator for constant-depth Frege systems. Constant-depth systems are relatively weak but not ridiculously so, and the list of proof complexity generators hard for them was until now rather short and did not include Nisan-Wigderson generators.

The proof of the main result once again takes advantage of a famous classical theorem: it exploits a notorious weakness of constant-depth systems, which is that they do not have short proofs of the pigeonhole principle. The symmetricity and hardness of f means that there is some *moderately* large u such that f has a different value for strings with u 1's than for strings with $u+1$ 1's, and the hardness of PHP means that, intuitively speaking, from the point of view of short constant-depth proofs there is no difference between u and numbers close to u. As a consequence, short constant-depth proofs are not able to tell that for a string with roughly u 1's the property f can have only one truth value, and in a sense (made precise by means of a nonstandard model of arithmetic) they cannot rule out not merely that all possible values of the the Nisan-Wigderson generator are obtained, but even that they are all obtained on the same input!

Paper C concerns jump operators in proof complexity, that is operations that take a (code of a) propositional proof system and output a strictly stronger one. Of course, the very existence of any such operator is open, as it would imply that there is no optimal proof system. The author proposes a new candidate jump operator based on combining Krajíček's concept of implicit proofs with interactive computation. Given a proof system P, the operator outputs a system $[\![\mathsf{IP}, P]\!]$ in which a proof has two main parts: the first is a succinct description of Prover's actions in a sum-check protocol intended to witness that the second part is a succinct description of a P-proof of the tautology at hand. (This is a probabilistic proof system, which will be a proof system in the sense of the orthodox Cook-Reckhow definition assuming standard complexity-theoretic hardness assumptions hold.)

Unlike in the other papers comprising the thesis, there is no apparent main result here. Among other things, the author illustrates the power of systems of the form $[\![\mathsf{IP}, P]\!]$ even for P as weak as treelike Resolution by showing that it is hard to prove lower bounds for them, in various senses. He also shows that if the soundness of $[\![\mathsf{IP}, \mathsf{treelike\ Res}]\!]$ is provable in a particular weak arithmetic theory, then the non-automatability of the Extended Frege system can be derived from hardness assumptions in structural complexity theory; currently, the non-automatability of EF is only known under cryptographic assumptions. Finally, there is an elegant contribution to the general theory of jump operators, connecting the existence of an effective jump operator to the conjecture that arithmetical theories T do not have short proofs of finite consistency statements for $T + \mathrm{Con}_T$.

The main aim of Paper D is to answer the following question: is every Kripke model of Heyting Arithmetic *locally PA*, i.e. does the classical model at any given node of a Kripke model of HA have to satisfy PA? This is a question dating back to the 1980's, and over the years there have been many partial results, mostly positive ones saying that every Kripke model of HA with a particular kind of frame has to be locally PA. However, the author shows that the answer is negative. In fact, he provides a characterization of the structures that can be at the root of a Kripke model of HA: they are exactly the structures satisfying the Π_2 consequences of PA; thus, not even Σ_1 collection is required.

The characterization is proved by a careful model construction, which relies on analysis of the properties of theories extending HA by the non-classical scheme ECT_0 (Extended Church's Thesis) and the diagram of a particular classical model. A crucial lemma says that whenever the classical model satisfies the Π_2 consequences of PA, such a theory has the existence and disjunction properties. The argument employs realizability in an essential way. A related second construction shows that models of HA that are not locally PA do not have to force highly non-classical axioms such as ECT_0. The proof of this is similar to the first one, but it additionally makes use of a formalized existence property and has a slightly metamathematical flavour.

Assessment of results. All the papers contain unquestionably novel results. Since none of the papers has a coauthor, the question of the extent of Mr Khaniki's contribution to the results does not arise.

In my opinion, papers A and B are best described as good (at a PhD level, perhaps even very good) but not spectacular additions to the proof complexity literature. Both of them contain one or two main results that provide meaningful new information about active research topics – the strength of Resolution over linear and polynomial equations in one case, the theory of proof complexity generators in the other. In paper A, for example, I would view the general size-width connection and the daglike $\mathsf{Res}(\mathsf{PC}_{d,\mathbb{F}})$ lower bounds as the main results. These are obviously interesting and relevant contributions, though it is hard to say whether they will turn out to be a helpful step on the way to results that would undeniably be major progress: in this case, say, superpolynomial or even superquasipolynomial lower bounds on daglike $\mathsf{Res}(\mathsf{PC}_{1,\mathbb{F}_2})$ and similar system.

In both papers, the proofs of the main results ultimately come down to a reduction to a major theorem from the past – the size-width connection for Resolution in paper A, lower bounds on proofs of PHP in B – which then does the heavy lifting. However, it should be emphasized that the arguments behind the reductions are clever and original, and they combine a number of ideas in a technically nontrivial manner. In paper B, for instance, the argument combines discrete probability-theoretic calculations with the use of a (previously known) characterization of the existence of small propositional proofs in terms of nonstandard models of arithmetic.

Paper C is relatively hard to assess. On the one hand, in this paper the author's quite impressive erudition is on full display: the arguments make use of concepts and results from a vast range of areas, from interactive proofs through formalized approximate counting to nonstandard models and Gödelian reasoning. On the other hand, there was no single result in the paper that stood out to me as particularly striking or fascinating. The theorems seemed generally interesting but had rather complicated statements and often included assumptions for which it was difficult to gauge their plausibility or relate them to more fundamental complexity-theoretic conjectures. The proofs were technically demanding, but many appeared to require the careful handling of sophisticated concepts more than genuinely deep originality. The result that probably appealed to me the most – presumably revealing my own limitations and biases – was the most "old-school" one, stating that the existence of a computable jump operator in proof complexity is equivalent to the nonexistence of small proofs of $\mathrm{Con}_n(T + \mathrm{Con}_T)$ in T for sufficiently strong theories T. This result also has a proof that is quite insightful and elegant, while at the same time not that difficult to follow.

The main result of paper D is clearly the gem of the thesis. Even though traditionally understood metamathematics of arithmetic is not quite as fashionable as some areas of complexity theory, it is very much a serious research field, and the author has solved a long-standing and extremely natural problem in it. This is an outstanding result, and I was also quite impressed by the way in which the problem is solved. The argument is original and ingenious but not overly complicated, and it blends a number of ideas (existence property, realizability, basic Kripke model theory; in the second construction also formalized provability) in a highly elegant way.

Throughout the thesis, I did not find any serious mathematical errors that would potentially invalidate a proof. There were some small mistakes that either had no influence on the flow of the argument or were easily fixable. In a few places, I had the impression that a piece of the argument was missing: two instances of this were Theorem 8.1, where it seemed that only one of the two parts was proved, and Theorem 15.6, where the proofs of both parts seemed to lack a final step. However, in each such situation it seemed clear that the missing part of the reasoning could easily be supplied and would go along predictable lines.

Presentation. As mentioned, the thesis has the form of a collection of essentially independent papers. I must admit that even though I understand some advantages of preparing a thesis in this form – in particular, less effort expended on some thankless tasks – I still have the belief that a traditional unified thesis is the preferable option. One benefit of traditional PhD theses is that they are often among the best sources for learning cutting-edge topics: a unified set of concepts and conventions, proofs that include details often omitted in research papers, and ideally an extensive introduction providing a thorough overview of the subject are immensely helpful in that respect.

With the partial exception of paper D, which has a more narrative approach, the papers share the same general structure. An introduction presents the results in some context but does not attempt to

define or explain all of the main concepts used. This is followed by preliminaries, precise statements of the main results, proofs of those results, and finally a concluding discussion that also mentions some open problems. Again, this has its advantages and drawbacks. An advantage is that the structure is clear and one knows roughly where to look for what, even if that takes a lot of page turning. The disadvantages become apparent when the results draw upon extensive background knowledge and involve many specialized concepts: then the reader will struggle to get a feel for the theorems if they are very far removed from their proofs. This is particularly evident in paper C, where the introduction makes use of (without explaining) such notions as 1-subEXP or exponentially pseudo-surjective, the preliminaries are very long, the section stating the results includes nine theorems, and the proofs begin well past the midpoint of the paper.

Fortunately, once the proofs do come, most of the time they are rather carefully written. In particular, the author makes an effort to present details of various inductive arguments, taking care to discuss even steps or cases that others would omit as routine. As mentioned, on a few occasions it looks as if a step in an argument is missing, but such situations are the exception rather than the rule. The amount of small mistakes in proofs and apparently missing or insufficiently careful explanations is the largest (by far) in paper C, though it can be observed that also paper B is not as clean in this regard as the two journal papers.

The amount of typos is generally acceptable and does not interfere with readability. Once more, paper C is to some extent an unfortunate exception, as the density of typos in it sometimes becomes annyoing. Still, that is at worst a minor distraction.

Overall, I would rate the presentation as reasonably good but somewhat uneven, with (unsurprisingly) the journal papers being the most polished and the unpublished paper C the least.

Conclusion. Erfan Khaniki's thesis contains a number of interesting new results, and it proves beyond doubt that the author is capable of carrying out original research work. Moreover, I have the impression that I can already see glimpses of Mr Khaniki's developing personal research style, which involves a significant dose of what I would call "creative erudition": the ability to turn very good knowledge of the literature in different research areas into a tool for uncovering new facts somewhere near the intersection of those areas. **I conclude that the thesis is clearly sufficient to serve as the basis for awarding a PhD degree**.

In my country, a doctoral degree can be awarded "with distinction". When treated seriously, this is an indication that the thesis might be worth considering for various national- and international-level awards. While reading Mr Khaniki's thesis, I asked myself whether I would vote for such a distinction in this case. I came to the firm conclusion that I would, largely on the strength of the paper on models of intuitionistic arithmetic. The three proof complexity papers would comprise a good thesis on their own, but it could then be criticized for the lack of a single spectacular result or for the uneven presentation. However, the theorem characterizing roots of Kripke models of HA more than makes up for those modest deficiencies.

Leszek Kołodziejczyk

Appendix. Detailed comments

[The notation n^m stands for m-th line from the top and n_m for m-th line from the bottom on page n.]

7^{12}: should it be $[\![\mathsf{IP}, Q]\!]$ rather than $[\![\mathsf{IP}, P]\!]$?

7_{21}: nonautomat*a*bility.

General comments on Paper A:

- There is something rather seriously wrong with the numbering of internal references in the paper (at least as included in the thesis). For example, there are references to Propositions 11 and 12, which should most likely be 5.1 and 5.2; on page 35, "Lemma 9.1" should probably be 5.3, and "Theorem 19.2" should be 3.2; a reference to "Section 4" at the bottom of page 25 should probably be to Section 5; and so on.
- In some parts of the paper (e.g. Theorems 2.1 and 4.1, but also elsewhere) it is assumed apparently without comment that n stands for the number of variables. But this is not a global assumption, as in some places it is stated explicitly (e.g. Theorem 3.2), and in other places it does not hold (most of Section 3.3).

Page 24, comment after Theorem 3.5: for unsatisfiability, clearly you need some lower bound on Δ as well? Perhaps the bound is the one appearing in Corollary 4.6?

Page 25, Lemma 3.8: something seems off, possibly you have switched from $\bigvee_i x_i$ to $\bigvee_i (x_i - 1)$ (i.e. from the logical to an algebraic understanding of the disjuncts) between the statement and the proof.

38^{10}: I don't see why you can have the -1 in the parenthesis: it looks like that assumes that the degree of $af + bg$ is exactly d and not lower. A similar thing happens in 40^8, but in both cases it should not have any significant effect on the overall bound.

Page 44, abstract of paper B, and again on page 47: in the list of applications of the theorem, the quantification "for any f" applies to both items 1. and 2., but it is confusingly located within item 1.

48^{10}: every f_n has a representation for any f whatsoever; the question is whether the representation has feasible size.

Page 49, Lemma 7.1: the lemma as such is a special case of the "obvious" fact that the truth value of formulas – of any sane logic – is preserved under isomorphisms (which are simply bijections if the language is pure equality). But in the application of the lemma in Section 9, we additionally need to know that the lemma remains true when statements like $\mathcal{A} \models \phi$ are evaluated within a model of the weak theory V^0_1. At that point, it begins to matter that ϕ is first-order.

50_9: "it" should be σ.

50_5: it would be better to say what "exponential" means here.

Page 51, Theorem 8.1: the theorem has two parts, but in Section 9 is looks as if only one of them is proved: apparently part 2. So, it should be explained what one should do to prove the other part as well.

52^{12}: the inequality $n^s 2^n \leq 2^{n^{t/u}}$ for an appropriate $u > \mathbb{N}$ should follow from the other assumptions: we are assuming $n^s < 2^n$, so $n^s 2^n \leq 2^{2n} \leq 2^{n^2}$, and we can take $u := t/2$.

54^3: g should probably be σ.

General comment on the proof of Theorem 9.1: it would have been very helpful to include an outline of the proof strategy, either on a general level or on a well-chosen simple example (for instance, $f =$ parity, A is the matrix in which in each row the 1's are the first n^t entries). That would let the reader understand that in the nonstandard model $(\mathcal{M}_{n^t}, \chi')$, the intuitive reason why everything is in the range of $\mathsf{NW}_{f,A}$ is that because of the failure of PHP, f is no longer an $\mathsf{NP} \cap \mathsf{coNP}$ function: e.g., the string α has a witness both for having an even and having an odd number of 1's.

Page 56, top half: I had problems with understanding the $\omega(\ldots, F(\ldots))$ notation, for example the third argument of F should be a number below n^t, so it was not obvious what strings like θ_0 were doing there. I guess the correct interpretation is that θ_0 should be treated as a *set* of some numbers below n^t, and $F(i, a, \cdot)$ should be understood as the image of $\cdot$ under F with i, a fixed?

56^{6}: misspelled "hardness".

56_{7}: misspelled "constructible".

56_{7}, 57^{2}: how should one understand the difference between "efficiently" and "effectively"?

57^{3}: should n^c be n^s?

57^{11}: generators *for*.

General comment on Paper C: there are many typos and minor grammatical issues (missing articles etc.) in this paper, too many to attempt putting together a comprehensive list. I only list the ones that might have a bearing on the mathematical content.

65_{15}: "They proved that (...) Ref_ϕ (...) does not have short resolution refutations". This requires ϕ to be unsatisfiable, and in fact Atserias-Müller did not prove that: they proved a width lower bound for Ref_ϕ and a length lower bound for a relativized version of it. The length lower bound for the unrelativized version was proved later by Garlík.

65_{2}: in $[\![\mathsf{IP}, P]\!]$, should the P have been Q?

66^{15-16}: for what proof systems are the assumptions true according to experts?

67^{12-10}: in your list of properties equivalent to the existence of an "efficient jump operator", how does the first item on the list differ from the statement it is supposed to be equivalent to?

Page 68, middle of page: the definition of Δ^b_n requires context – is the equivalence provable in a theory, true in a model, locally assumed as part of some argument? – and this should be pointed out because this context may vary, and various statements will be true or not depending on what it is. Cf. the usage around Theorem 12.1 with the one at the beginning of Section 12.7 – in the latter case, you have provably Δ^b_1 properties in mind.

70^{5}: I could not understand what $Count(C, y)$ does and why it exists. In $|X_C \cap y|$, if the vertical lines mean cardinality, then you probably need to assume $y \in Log$ already here, like you do with 2^k in the next sentence? Or do the vertical lines mean length? But then I do not understand what this notation is saying.

Page 71, item 2. of Theorem 12.5: x_C should be capitalized.

Page 71, Definition 12.3: the empty square and empty circle should probably be the same symbol.

Page 75, around Theorem 12.11: when you say that we get the theorem "from the above definition", do you mean that it follows immediately from the definition? If so, then I do not see it.

Page 75, Theorem 12.13: is there a specific reference where the theorem is stated in such generality? E.g. a chapter in one of Krajíček's books?

Page 77, Definition 12.20: what is c_0 supposed to do? All of a sudden it appears at the bottom of the page, but it plays no role in the definition.

Page 78, Definition 12.23: this definition is not particularly elegant (it looks like it redefines a concept already defined in Definition 12.22). Also, it should once again explicitly specify what the difference between $\mathsf{P/poly}$- and $\mathsf{NP/poly}$-naturality is.

79_{6}: it should be explained what arithmetizing means in this context.

79_{1}: is the prover really required to send the specific polynomial Q_i? My understanding is that the prover is allowed to cheat, and that the whole question is what happens if he does.

Page 79, Definition 13.2: you want the sentence ϕ to be true.

80^{2}: there is an $\in$ sign missing. Also, you do not say what happens if the sum-check fails.

Page 80, middle of page. What do you mean by (the truth-table of) π being the transcript of the prover? Should it be a description of what prover does for all possible r's? It is hard to "note" a bound on the size of π until it is known what π is supposed to be.

Page 80, Theorem 14.1: you should indicate that you are going to prove (a formalized version of) the theorem later on: otherwise a reference would be expected.

Page 80, Definition 14.1 and below: note that in the definition of MA proof system, you did not use the term "verifier" that is used here. It takes a while to grasp that Definition 14.1 intends to define $[\![\mathsf{IP}, P]\!]$

directly as an MA proof system, while "IP-randomized implicit proof system based on P" is part of the term being defined, not of the condition defining it.

Page 81, Definition 15.1: has the notation $\preceq^f_\epsilon$ with superscript f rather than α been introduced? Anyway, I would have thought that f is a boolean function given as a truth table, in which case its size should be a power of 2 rather than the value given in item 1.

82_{15}: should the subscript Q in the exponent under the probability symbol be P?

Pages 82-83, Theorem 15.8: this theorem is virtually unreadable without a review of what *all* the constants involved mean. It would be strongly advisable to provide such a review, either just before the theorem or elsewhere in a separate table with references to where the constants first appear.

Page 83, Theorem 15.6: in both items, in part (a) you might want $\mathcal{M}^*$ rather than $\mathcal{M}$. More fundamentally, you use the notion of $[\![\mathsf{IP}, \mathsf{Res}^*]\!]$-proof, so you should say explicitly what a proof is in an MA proof system. If I am guessing correctly, that notion is probabilistic, so it should be said why the relevant probabilistic statements make sense in $\mathcal{M}^*$ even though the model is not assumed to satisfy any weak pigeonhole principle.

86^{11}: why can we assume that ϕ has *exactly* 3 numerals per clause? This has a slight influence on the bound on m used later on.

86^{17}: why is p in $LogLog$ rather than just Log?

Page 86-87, description of formula $\psi(k)$ and later on. The formula ψ and its inductive proof should be described better: e.g. what is v and how exactly are the polynomials P_r related to the polynomials Q_i discussed on pages 79-80? Also, I think that many indices are "off by 1": the induction starts at $k = 1$, but that would correspond to $i = n$ (while you apparently want $i + 1 = n$), and it ends at $k = n - 1$, but that corresponds to $i = 2$. Finally, I think the explanation for the bound $\frac{m}{p}$ in the case where the verifier chooses a point from S is strange: isn't the point simply that the probability of choosing such a point is at most $\frac{m}{p}$?

88^1: how does "the power of 1-EXP" make F' Δ^b_1-definable? Is the point that it can be Δ^b_1-defined using a large enough additional parameter, which the exponential overhead will let us provide?

88^8: constant, not contact.

Page 89, proof of Theorem 15.1. I got confused as to what counts as a "short" T-proof here. Does "short" mean of size polynomial in $2^{\max(n,m)}$?

90_6: since it is important that $2^{t^{c_r}+c_r} \in Log$, an explanation why this is the case, with a reminder of how d, d_Q, c_r are involved in defining t, would be in order.

92^{5-7}: I do not think that the statement of Theorem 15.2 as such implies the probability bound you want here, because in $Sound_c$ the probability of erroneous acceptance is only approximately bounded by 1/2 with approximation error $\epsilon = 1/4$. To get the bound, you need to invoke the argument proving Theorem 15.2.

Page 92, Theorem 16.2: do you want "*for all* $k \in \mathbb{N}$" or perhaps simply "let $k \in \mathbb{N}$ and suppose that..."? The existence of k for which $tt_{n^k,n}$ is not hard comes down to $k = 1$ and may well be provable (by verifying some $(1 + \epsilon)n$ lower bound for an explicit function in S^1_2).

92_7: you probably want p, u, f to depend only on n, not on φ.

93^{12}: "true" should be "tautological".

Proof of Theorem 15.6: to complete the proof, you should say how $\mathsf{NP} \subseteq \mathsf{P/poly}$ implies the existence of a $\mathsf{P/poly}$ natural property against $\mathsf{Size}(n^k)$ (and similarly for the second item of the theorem).

95_9: "the assumption of the theorem" should be "item 2. of the theorem" or "our assumption that item 2. holds".

Section 16.9: instead of expressions like "polynomial size proofs of $A(n)$ in n (where A is some statement), it would be better to use "proofs of $A(n)$ of size polynomial in n".

Page 96, Lemma 16.4: you probably wanted to say that Φ_F defines the graph of F in the standard model.

97^{10}: the sentence ϕ should be $\forall\Pi^b_1$.

97_{18-17}: why not say explicitly what η is?

98^{6}: this formula might not be entirely correct because some fixed size is needed to prove T from S; thus x might be offset by an additive constant.

110_{9}: construction*s*.

112^{12}: $\psi(x)$ should be $\phi(x)$.

Proof of Theorem 20.5: I was slightly confused by the lack of a discussion of how the tuples $\vec{c}_i$ overlap with $\vec{d}$ and with each other, and which parts of the tuples the variables $\vec{x}_i$ correspond to. A similar comment applies to Lemma 20.12.

118_{13}: unneeded "Note that".

119_{15}: "this version" should be "these versions".

120_{10}: you should point out that in the definition of $\mathsf{EXT}(\mathcal{M}, T)$, the role of constants for the elements of M is played by the formalized numerals, which is important if $\mathcal{M}$ is to believe that T proves anything substantial about them.

122^{18}: well-know*n*.

$122_{1}2$: n is used in two different roles.

122_{9}: the negation between $\exists x$ and ϕ should be deleted.

125_{3}: "following the same argument". It was not clear to me which argument is meant.

126^{7}: unneeded "So".

126^{13-14}: "this shows the known positive results are the best we can get for BA and EBA". I did not understand this. For example, for EBA the argument based on HA gives the Π_2 consequences of PA as an upper bound, while the positive result you mention is about a weaker theory.