The management and exploitation of phosphorus in bloomery iron: A practitioner's observations

Lee Sauder and Skip Williams

ABSTRACT: Phosphorus was a very common alloying element in ancient iron. Though phosphoric iron has been increasingly studied by archaeometallurgists in recent decades, little work has been published on how preindustrial smelters and smiths may have controlled and utilized phosphorus in their iron. This paper briefly summarizes our experience in controlling the phosphorus content of bloom iron. Part one deals with controlling phosphorus levels through various smelting parameters, including the speed, temperature, and length of the smelt, the proper management of slag, and the addition of calcium. Part two discusses the removal of phosphorus after smelting. Part three presents some observations on the positive aspects of working with phosphoric bloom iron.

Introduction

First, a word to help you understand the perspective from which this paper is written. I am not an archaeologist, a scientist, or a scholar. I am a professional artist-blacksmith with 51 years of experience. Since 1998, reviving the nearly lost techniques of bloomery smelting has been the dominant focus of my work. My friend Skip Williams has been my primary collaborator in puzzling out a few of the mysteries of this craft. As of this writing, I have conducted almost 300 smelts, producing more than 2 tons of bloom iron. I have forged almost all of that iron to finished products, mostly sculpture, but also utilitarian things like knives, axes, and hammers.

I have traveled internationally to consult with archaeologists, and to teach and demonstrate this craft, using many different ores and many different furnace designs. So my experience in working with bloom iron is both deep and broad. But I do not have much experience in, nor much access to, modern analytical methods. So given the unusual nature of my experience, and my limitations, perhaps you will understand and forgive this paper's deviation from the usual style and standards of this journal.

In the fall of 2001, an alert and peripatetic friend told me of a pile of nice looking ore he had found during his Virginia mountain ramblings. At that time, our best ore source was in an abandoned, partly collapsed underground mine, and gathering ore there was becoming increasingly nerve-wracking. So I immediately went to investigate this new source, and was thrilled to find a small surface cut with the extracted ore piled neatly alongside. The ore was a dark, dense and rich limonite/goethite. And there was lots of it!

When this mine had been abandoned, most likely in the late 1840's, the miners did not know they were not returning the following season. So the pile was the actual mine run, rather than the tailings we had been picking through at other surface mines. Being rich, plentiful, and reasonably accessible, this mine provided the ore for most of my smelting experiments over the next decade.

ISSN 0142-3304 (print)
ISSN 2755-0249 (online)

By summer of 2005, it had slowly, and sometimes painfully, dawned on Skip and me that we were dealing with a rather highly phosphoric ore. Much of our ensuing experimentation, and the bulk of our pondering, from 2006 to 2012, concentrated on learning to manage the phosphorus content in our iron.

Some of what we learned has come through deliberate experiment, some through serendipity, and some through the slow accrual of the vague and half-forgotten impressions that occasionally fall into recognizable patterns. My understanding of these issues still feels rather loose and rattling, but I will do my best here to share what I think, despite my feeling of peering at dim shapes in an early morning fog. In part one, I will discuss some methods for controlling phosphorus during smelting. In part two, I will describe a method to remove phosphorus from bloom iron. And in part three, I will talk about the pleasures and challenges of working with phosphoric bloom iron.

Analytical methods and shop tests for phosphorus

Before I describe our methods for managing phosphorus, I think it is important to describe the methods we used to measure that phosphorus, as they are not (yet) conventional. But they have proven quite useful. The more accurate and quantitative laboratory analyses that is daily bread of this journal's readers were of very limited availability to us as we explored these questions, and so we had to develop some more immediate and tactile tests that would help us to judge the results of our experiments. We do have a little metallographic lab set up, and we used this to help verify the efficacy of our shop tests. Our metallographic skill is not such that we could use it to quantify specific amounts of phosphorus, only relative amounts. We also corroborated some of our twist tests using Klemm's reagent to see higher and lower phosphorus areas macroscopically. We did comparative hardness testing. And finally, we did have a few lab analyses done later that also confirmed the value of these shop tests.

I think modern analytical methods are a bit problematic in this situation anyway, because they are actually too precise. In a material as heterogeneous as bloom iron, the condition you are observing may be quite different from what sits a few millimeters down the bar. This can make any quantitative analysis from a small sample set quite misleading. Even metallography only reveals the chosen surface plane. Also, most quantitative analyses cannot easily tell you whether the phosphorus is present in the metal or the slag. Besides, knowing the actual phosphorus content of the iron would be useless information to an ancient smith. And if we are trying to recover their lost knowledge, it is well to keep in mind what mattered to them. It is the working properties of the metal the smith is concerned with, and the iron's suitability for the job that mattered to the user. So we concentrated on developing methods of shop testing to monitor our phosphorus content, and tended to use various lab analyses, metallography, and hardness testing as secondary information.

Also, in trying to continuously improve my bloom smelting, I found it is very important to work some of the bloom into bar immediately after smelting, so that the subjective experience of the forging is directly connected to the experience and memory of the smelt. Lab analyses that happen days or months later cannot help you rapidly move forward in this way. One of the classic shop tests for iron is to notch the metal, bend it until it breaks, and observe the fracture. We found this useful to begin with, but it suffers from the same problem of excessive precision: it only shows you the condition of the iron at the single point that you have chosen to notch. Nor is it very quantifiable. We tried to measure the angle of bend before failure, but we found this rather difficult to do with any accuracy.

The most useful test we have devised, by far, is twisting the metal until it fails and breaks. This twist testing has several advantages. First, the twisting reveals the character of the bar over the length that you choose to twist, giving a bigger window into the bar than single cross section of the notch test, or the single plane surface chosen for metallography or other lab testing. Second, it is easily quantifiable. By counting the number of revolutions before failure, you have a simple way to record your result, and compare it to other bars, past or future. Third, the act of twisting has a direct connection to the hand that gives a valuable subjective impression: stiff or ductile, smooth or crunchy. And twisting happens more slowly than breaking, giving time to observe the bar under stress, especially what happens on the corners of the bar. And then finally, you still have the visible fracture surface as you would with the notch and break test.

We devised two twisting tests that we find useful, one, a test of the metal while hot, the other a test while cold. I would propose these, especially the cold twist test, as a standard that other bloom smiths could follow, so we can compare materials more easily. A more complete and specific description of these tests follow in the appendix. Here, just a few words of description will suffice. These tests are based on the well-known attributes of high-phosphorus iron: ductility at high temperatures, brittleness at cold temperatures, and unrestrained grain growth at high levels of phosphorus. In the hot twist test, the higher-phosphorus bars will survive for more twists than the low-phosphorus bars will, as the high-phosphorus metal is more ductile at forging temperatures. In the cold twist test, which is our most helpful and accurate test, the high-phosphorus bars snap at less than half a twist. The low-phosphorus bars survive from one to two full revolutions before breaking.

The cold twist test has proven so useful that I almost always forge a bit of bloom to bar and test it this way, immediately after the bloom comes out of the smelting furnace. Metallographic examination, etching with Klemm's reagent, and a few lab analyses for phosphorus have confirmed the reasonable accuracy and usefulness of this test for judging phosphorus content.

Part one:
The management of phosphorus in the bloomery

Overview of experiments

Of course, in my first few decades of smithing with modern steels, I had never encountered phosphorus, since it has been as carefully eradicated from our steels as if it were smallpox. And since bloom iron is so variable anyway, and I did not then know anyone else with significant experience in forging bloom iron with whom I could compare notes, it took me quite a while to realize that we were dealing with phosphorus in our iron. The first attributes of phosphorus I noted were the positive ones: a wonderful softness and ductility under the hammer; ease of forge welding; a beautiful banded sheen when polished; a resistance to corrosion.

The focus of our experiments at the time was on increasing our efficiency in time and labor, and increasing our yield from the ore. Unbeknownst to us, many of the methods we used to increase yield also tended to increase the phosphorus content of our metal. Over time, here and there, the negative aspects of phosphorus began to appear: the occasional breakage of an attenuated form when I tried to adjust it cold, and many crazy interactions with carbon when we smelted steely blooms, or tried to carburize our iron after smelting. By the time a knowledgeable archaeometallurgist, Evelyne Godfrey, analyzed a particularly troublesome bit of our iron, and told us it had about 1 % phosphorus, we were not too surprised!

And so we began to try to understand how phosphorus behaved in bloom iron, and how to control it. Here our partnership fell into its common productive pattern: I made iron and banged on it, watching and feeling what happened, while Skip puzzled through the available literature, excavated the salient facts, and helped me understand them. Some of the experiments were conducted in planned campaigns specific to the phosphoric questions, and some were dispersed through the years as I just made things. I will not attempt to recount all the experimental data and results here, which would fill this entire volume. I will simply report our current working understanding that arises from those experiments.

Smelting procedures that affect phosphorus content

Of course, the simplest and most obvious way to control the chemistry of the iron that comes from the furnace is to control the chemistry of what you put into it. But the few ore analyses we managed to have done do indicate that the phosphorus can vary extremely from rock to rock, and strata to strata within the rock. There seems to be no reliable visual indicator of which ore is high in phosphorus and which is not. This hidden variable causes a huge degree of uncertainty (which is one reason we needed 60 or 70 smelts to even begin to get a handle on this). For each experimental campaign that focused primarily on phosphorus, I tried to minimize this variable by roasting, and breaking large quantities of ore at one time, enough for six or eight smelts, and then mixing it well. So if we cannot control the phosphorus content of our ore by sorting it, the next obvious question is: can we influence the final phosphorus content of our metal by altering our smelting procedure?

The furnace design and general smelting method we used to answer this question is covered rather comprehensively elsewhere (Stepanov et al. 2022), so I will only give a brief description of the furnace here. The furnace was a clay shaft of 25 to 26 cm inner diameter, about 1 meter tall, tapping the slag outside the furnace, rather than into a pit (see Fig. 1). The blast was supplied by an electric blower through a copper tuyère with an orifice of 2.2 to 2.3 cm. The bloom forms well below the tuyère, filling the diameter of the furnace, and the slag is tapped from below the bloom.

I have been able to identify several smelting variables that influence phosphorus content, and they have all been arrived at empirically. Only slowly has a hypothetical framework emerged to start to tie them together. So let us just deal with the variables first, and then the hypotheses later, just as we had to do in actuality. The variables

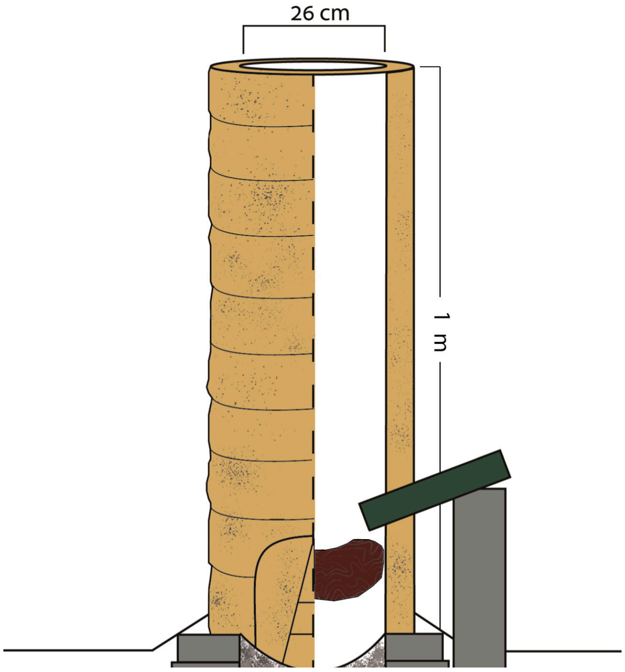

Figure 1: Schematic diagram of furnace showing bloom position. (Fig. 2a in Stepanov et al. 2022, licensed under CC-BY 4.0).

affecting phosphorus content that I have (so far) been able to identify with some degree of confidence are:

1. The tempo of the smelt
2. The slag tapping procedure
3. The charging procedure
4. The overall length of the smelt, and
5. The addition of calcium as flux.

Though these variables are most effective taken in concert, for the sake of clarity I will discuss them one by one, roughly in order of importance.

Smelting tempo

By tempo, I am referring to the interrelated variables of air rate, ore:fuel ratio, and temperature. I will flesh this out a bit, but here is the takeaway: smelting too hot and fast tends to increase phosphorus content. Smelting too cool and slow inhibits phosphorus uptake, but also tends to raise carbon content, and depletes the iron in the slag, making for poorly working iron. So you have to hit the sweet spot in the middle.

In my usual smelting practice, one of my main controls is the elapsed time to consume a 1.8 kg (4 lb) bucket of charcoal, along with its attendant ore charge, which I will refer to as the charge time. We can alter the charge time either by changing the air rate (affecting both temperature and the rate of descent of the ore), or by altering ore:fuel ratio (affecting both temperature, amount of reduction, and thus slag chemistry). When smelting for best results rather than experiment, I keep these variables in a rather narrow range, with an ore:fuel ratio near 1:1 by weight, and with an air rate that consumes 0.35 grams of charcoal/minute/cm^2 of furnace cross-section (Sauder 2013a, 72).

To illustrate the delicacy of this tempo, in the series of experiments in which we first identified this variable, in a 25.4 cm (10 inch) interior diameter furnace, a charge time of eight minutes yielded high-phosphorus metal. A charge time of 12 minutes yielded a low-phosphorus but high-carbon iron, that was difficult to forge into a bar, and difficult to weld, with transverse cracking. A charge time of 10 minutes yielded a low-carbon, lower-phosphorus, tenacious iron that forged into bar quickly and easily, with no cracking. In short, the proper tempo is not adagio, not presto, but a nice comfortable andante.

Slag tapping procedure

When and how you tap the slag from your bloomery has a major effect on many aspects of bloom quality: its density, its carbon content, even its shape. And yes, its phosphorus content. In our earlier practice, when we tended to be more concerned with yield than quality, we generally waited to tap slag until near the end of the ore charging. We then charged the tapped slag back into the top of the furnace for an hour or two to reduce it further, then perhaps added another charge of ore at the end of the smelt to help decarburize the bloom. During this period of slag recycling, we would watch the character of the slag change as its iron content dropped, becoming more glassy, more viscous, and more black, rather than the grey wüstite color and metallic sheen of the earlier slag.

In the fall of 2010, I had the good fortune to spend a week helping to excavate a medieval bloomery site in Wiltshire, UK. Before this (other than a brief overwhelming afternoon at a 12th-century smelting site in Burkina Faso), almost all the bloomery slag I had seen was made by me. This was the first time I had an extended time to study the slags made by a past master of the craft. I was struck by how consistent it all was, and how it looked like the slag I tapped early in the smelt, before we cooked extra iron out of it by recycling it. This suggested to me that I needed to tap early and often, and not recycle the slag. When I returned home and tried this, I made a bloom that forged more sweetly than any iron I had ever made.

As I continued to refine my slag tapping process, I kept heading more and more in the direction of ensuring a constant flow of fresh slag, of high iron oxide content, through the incipient bloom. This constant and diligent slag tapping has several beneficial effects, but among those, it does indeed help to keep phosphorus content low.

Deciding when to commence slag tapping is a crucial decision. If you tap too soon, you will disrupt the slag bowl and incipient bloom, and never give the bloom a chance to form properly. Tap too late, and the iron will already be depleted from the slag, which will then not easily flow out of the incipient bloom. As I tried to figure this out, I usually started tapping after about 50 or 60 percent of my intended charge was in the furnace. But over time I noticed a clear cue that helps me see when it is safe to begin tapping. In the earlier phase of the smelt, there will be copious little jets of flame coming out of the little cracks and edges of the tap arch. But as the slag bowl, and then the incipient bloom forms, these flames lose pressure and then finally stop, as the slag bowl and bloom seal off the bottom from the blast pressure above it. Then you know that the bloom is sturdy enough to begin slag tapping. I then keep the slag flowing until the end of the smelt (see “Length of smelt” below).

When the slag bowl has formed across the diameter of the furnace, and is stable, you can open the tap arch, and clear a space below the bloom without it collapsing to the bottom of the furnace. From this point onwards, the tap arch remains open, so I can see whether the slag is flowing, and where the slag is flowing from. I tap the slag by scratching or gently poking the bottom of the slag bowl in the center of the bloom and furnace, and the slag will flow freely for a while, and then slowly freeze up, and then I can scratch and poke the bloom bottom again to renew the flow.

If you are having trouble visualizing this, think of a drip coffeemaker, with the bloom being the ground coffee, and the slag being the coffee dripping out of it. By having the slag flow out of the lowest part of the bloom, I can ensure that fresh slag is flowing throughout the bloom. If you tap above the bloom, or tap at the edge of the bloom so that the slag is flowing over the rim of the bloom, you are in danger of letting the slag within the bloom stagnate, and become depleted of FeO. This not only leads to more reduction of phosphorus, but leads to a more refractory slag. One should always remember that wrought iron is a composite material of iron and slag, and so the softness and fluidity of the slag at forging temperature is an important part of the quality of the bloom.

Charging procedure

As mentioned above, we used to continue the blast for as much as an hour and a half after the last ore charge, recycling tapped slag. This procedure was great for wringing the most iron from a given weight of ore, but it did tend to produce higher phosphorus iron. For a lower phosphorus content, it is better to continue your ore charges right up until the bloom is removed. Ore to fuel ratio will also have an effect. Higher ore:fuel ratios tend to produce lower phosphorus iron. I usually charge 1.8 kg (4 lbs) of ore to 1.8 kg (4 lbs) of charcoal, and push that up to 2.3 kg (5 lbs) ore to 1.8 kg (4 lbs) charcoal in the latter stages of the smelt. This saves fuel and time, as well as inhibiting phosphorus uptake.

Length of smelt

The longer the smelt, given the same other parameters, the higher the phosphorus content tends to be. I think a lot of phosphorus uptake often happens in the latter stages of the smelt. As the bloom fills in and becomes more solid, it can entrap slag, thus leading to more reduction of phosphorus from the slag to the metal. So if you want low-phosphorus iron, the bloom should be removed as soon as you think it is solid.

In general, not only in regards to phosphorus, but also in regards to carbon content and slag chemistry as well, after the bloom becomes solid enough that the slag will no longer drip through the bloom, the quality of the bloom will start to suffer, even as the total yield of iron is increasing. If you get too greedy about the yield, you will not only lose quality, but spend way more time forging the bloom into an adequate bar, as it will require more folding and welding, leading to more forging losses. So your perceived efficiency in smelting can end up being a loss of material and an increase in labor by the time you have actually made a useful bar of iron.

Addition of calcium

This seems to work to inhibit phosphorus in the iron, but it is not a method I typically use, or would recommend to others. We have found that adding calcium to the ore charge, in the form of ground oyster shell ($CaCO_3$) does tend to lower phosphorus content. Unfortunately, it has other negative effects that far outweigh its ability to lower phosphorus. Too much calcium leads to a very viscous, iron-poor slag, and it also encourages carburization of the bloom, both of which lead to recalcitrant, troublesome iron. Because a slag with too much calcium is worse than a bloom with too much phosphorus, I generally avoid this practice. The amount of calcium present in the fuel ash is plenty, the addition of more calcium is usually unwise.

Discussion

So I have set out some steps that we can take to influence the phosphorus content of our bloom iron. Some so simple, and subtle, that they could be easily overlooked by someone who is new to the game, or even an experienced smelter who is just out of practice. But, of course, this leads to the question of how these methods influence phosphorus content. This is important not just to satisfy our curiosity. Some sort of theoretical framework arising from this experience helps to guide our choices as we work, and if that theory keeps leading to better results, then it is a good theory! I guess it would be better to say this is a set of hypotheses, rather than a theory. Nevertheless, what follows is the way I currently perceive the behavior of phosphorus in the bloomery. I make no claim here that these are all settled questions.

Our usual goal when smelting with a phosphoric ore is to prevent too much phosphorus from reducing and alloying with our metal. If we fail to do so, it is also possible to remove phosphorus that has already reduced (which we will explore in part two). Both processes may be happening within the bloomery, I do not know. But as I make my choices during the smelt, I am simply thinking in terms of prevention.

From my reading on the subject, it seems to me that the actual mechanisms and circumstances that lead to the reduction and alloying of phosphorus into iron are rather complex, unpredictable, and not terribly well understood. (Godfrey 2007, 77-90; Vizcaino et al. 1997). It has been observed in archaeological iron that phosphorus is depleted adjacent to slag inclusions (Dillmann and Balasubramaniam 2001), but I have not seen the mechanism for how this happens explained anywhere. There are two basic ideas that guide my choices, when trying to control phosphorus in the smelt. I am not aware of actual science to back these ideas up. These ideas are just a framework for my decision making that have come from observation and comparison over these many smelts. First, it seems to me that the iron reduces just a little more easily than the phosphorus does. And second, it seems to me that the phosphorus requires a bit higher temperature to reduce than the iron does.

It follows from these two ideas, then, that if we do not want phosphorus, we simply have to make sure that there is always lots of iron oxide available for reduction, and that the furnace does not get way too hot. Our furnace will not reduce much phosphorus if there is always FeO available to reduce first, and the furnace stays at a moderate temperature. Another way of expressing that we want to have plentiful FeO available, is to say that we want to ensure that we do not get too much reduction. If we look at it this way, we can see how all the procedures above could work to limit phosphorus uptake.

Having a high ore:fuel ratio will limit reduction, as there is less carbon available for reduction. A smelt of an overall shorter time will obviously lessen the amount of reduction that happens. And by frequent tapping, we can get the slag out before it has been reduced too much. The question of the smelt tempo is a bit more complicated, but fits into the same scenario. The effect of the air rate is not a nice simple issue. If we turn up the air, we are increasing the temperature (thereby increasing the rate of reduction), but we are also speeding it through the furnace faster (thereby lessening reduction). The balance between these effects is not a nice straight line: that is, one does not offset the other at the same rate.

So this explains why our best iron results at an andante tempo. Too slow, and the ore travels through the furnace too slowly, reducing too much, and thus picking up too much carbon. But it is also at a low enough temperature that the phosphorus does not reduce. Too fast, and though the iron travels through the furnace faster, lessening the residence time, it is also at a higher temperature, increasing reduction. And the temperature at the hot zone increases enough that the phosphorus reduces more readily. As the above illustrates, all the procedural changes are affecting the others. Like everything else in bloomery smelting, there is no way to truly change one variable at a time. All are interrelated, and not always in nice clear simple ways, either.

Finally, we come to the addition of calcium as a flux. It is well established that phosphorus will bind preferentially with the calcium, thus keeping it in the slag (Kumar and Balasubramaniam 2002). Now, how the use of these variables might show up in the archaeology is, of course, outside my area of expertise. But there is one thing I have noticed in reading so many publications on the archaeology of smelting sites. It is very common for archaeologists to see a high FeO content in a slag, and say that this represents an inefficient smelting process. But I think a slag high in FeO is just as likely to represent a desire on the part of the smelter to limit carbon content, phosphorus content, or both. Whether for an archaeologist or a practitioner, a focus on yield of bloom from the ore is short-sighted, and does not take into account the ease of bloom smithing, the quality of the final material, or the type of iron required for its intended use.

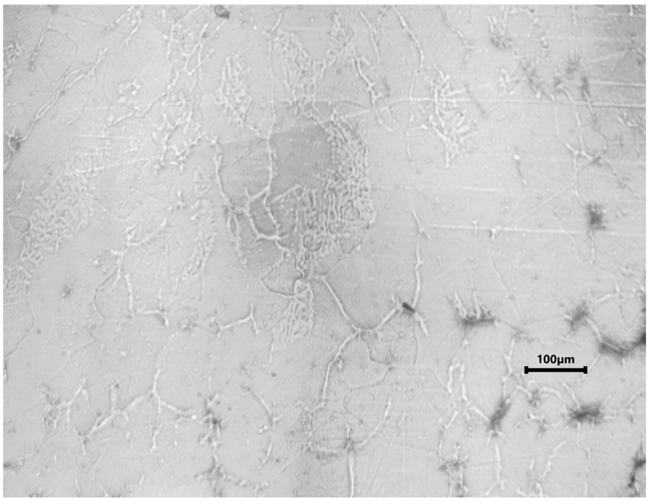

Figure 2: Smelt #177, smelted by high phosphorus procedure.

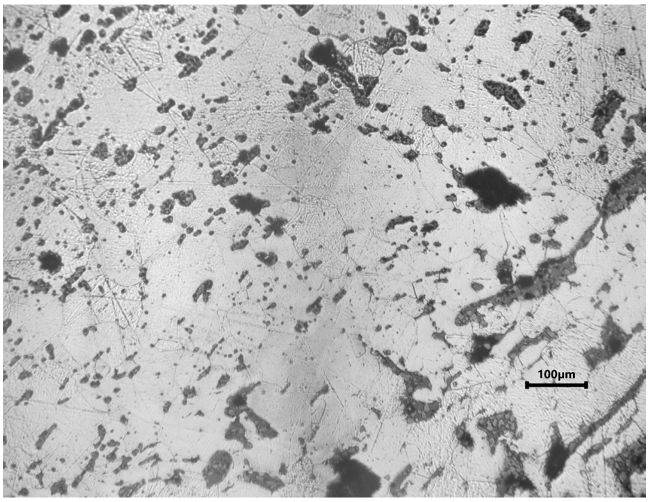

Figure 3: Smelt #173, smelted by low phosphorus procedure.

Analytical results of phosphorus campaigns

Our first campaign of smelts to explore how to influence the phosphorus in bloom smelting was comprised of 6 smelts in March of 2010. Here, we were just trying different things and seeing what happened. It was during this campaign that we settled on the twist tests as good indicators of phosphorus content. This set of experiments gave us out first glimmer of the smelting techniques outlined above. In late 2012 a colleague asked for a set of samples for XRF analysis, for the purpose of answering a different question: whether the use of the copper tuyère is detectable in the resulting bloom iron (it is not). Though we needed a range of samples with different ores to answer this question, I was able to include two samples from the 2010 phosphorus experiments. An earlier analysis of this ore (from the same mine, but not a sample from the actual experiments) showed a phosphorus content of 2.4 % phosphorus.

These analyses of samples from fully worked bars were done on a Bruker WDXRF Series 8 Tiger. These three analyses are the only quantitative measures of phosphorus levels that we have. I know it is a tiny sample set, but it does offer some confirmation of both the efficacy of these procedures, and the relation of the twist test to phosphorous content. Smelt 114 was run hot and fast, with a seven to eight minute charge time. Slag was tapped for one and a half hours and was recycled. The analysis of the iron showed a phosphorus content in the iron of 0.4 %. The cold twist test went half a twist before failure, and showed big grains on the crystalline fracture.

Smelt 129, from the same batch of the same ore, was with a moderate nine to 10 minute charge time, better slag tapping technique, and no recycling. The iron had a phosphorus content of 0.15 %. The cold twist test went one and a quarter before failure, and showed a ductile fracture with no grains visible to the naked eye. This is indeed a tiny data set, but at least it does tend to confirm that the twist test is related to the measured phosphorus content. After the 2010 phosphorus experiments, I continued to refine my smelting methods, especially as I began to sell bloom iron to other blacksmiths, as well as for my own use, so I had a strong incentive to keep getting better. This is the period in which I developed the low-phosphorus smelting procedures I have outlined. This period comprised of 60 smelts.

In 2014, I did another series of dedicated phosphorus experiments to confirm the utility of these methods. I prepared enough high-phosphorus ore for six smelts, and then alternated smelting for a high-phosphorus result, then for a low-phosphorus result in the next smelt. Each bloom was quartered, and each quarter was forged to bar, and the cold twist test performed on each bar. The results confirmed what I thought I had learned.

The high-phosphorus procedure yielded larger blooms from the same amount of ore, on average 7 kg. The cold twist tests ranged from a quarter to one twist before failure, with an average of 0.7 twist. The low-phosphorus procedure, intentionally less reducing, yielded smaller blooms, on average 5.2 kg. The cold twist tests ranged from half to one and a quarter twists before failure, with an average of one full twist. I had no opportunity for quantitative lab analyses on this series. I thought I had long since used all this iron, but while writing this paper I found two securely labeled bars from the 2014 experiments, and examined them metallographically. Figure 2 shows brittle iron from smelt #177, smelted by the high-phosphorus procedure, with the ghost structures characteristic of phosphoric iron, and a low slag content.

Figure 3 shows tougher iron from smelt #173, smelted by the low phosphorus procedures, in which we could not locate any ghost structures. Note also the plentiful and well dispersed slag stringers resulting from this method.

But the most important result was more general and subjective. The experiments produced two very separate and distinct characters of iron. The higher-phosphorus procedure yielding iron so brittle as to be useless for many conventional purposes, and the lower- phosphorus procedure yielding a nice tough, useful iron.

Summary

So, to review and summarize our smelting procedures. If we want to limit the amount of phosphorus in our iron, we will smelt at a moderate tempo and temperature. We will smelt at a fuel:ore ratio of 1:1, and increase the relative amount of ore if possible. We will tap slag as early as we safely can, and tap often thereafter. And we will stop smelting as soon as we have a bloom of adequate density and size, rather than getting greedy for more weight. We may add small amounts of calcium to our charges. Conversely, if we want to increase the phosphorus in our iron, we will smelt hotter and faster. We will smelt with a ore:fuel ratio of 1:1, or with slightly more charcoal. We will not tap slag as early or as often. And we will smelt longer, perhaps recycling iron-rich slag, and make a bigger bloom.

Part two: Reducing phosphorous after smelting

But suppose we end up with iron of an excessive phosphorus content anyway, either because we had a very high-phosphorus ore, or we did not control our smelt as well as we should have. All is not lost. The phosphorus can be removed rather easily. By remelting our bloom in a shallow hearth, and fluxing it liberally with iron rich slag, we can move the phosphorus from the metal into the slag. I have written elsewhere about Ole Evenstad's 18th century descriptions of converting iron into steel (Wagner 1990; Sauder 2013b). But I do not think it is widely understood that the first part of Evenstad's two-step process is for the purpose of removing phosphorus from the iron before converting it to steel. When I first read it, I assumed it was a method for consolidating poor blooms, and I think this is how others have generally interpreted it too (Vike 2016; Wagner 1990). But as I became more aware of the phosphorus in my ore and my iron, I began to wonder if that first step was instead a method of removing phosphorus, which is such a common component in the bog and lake ores Evenstad was smelting.

In March of 2011, I convened a crew of friends for ten days of experiments, including a quite comprehensive series of Evenstad remelts, and found that the first step of remelting did indeed lessen the phosphorus content, at least according to our twist tests. In the same set of XRF analyses I described in the previous section, I included one sample of phosphoric bloom, and one sample of that iron remelted by Evenstad's process. This showed the phosphorus decreased from 0.42 % to 0.36 %. My friend Tim Young, who was present for this experimental series, later did an SEM analysis of iron before and after remelting. He saw that the overall phosphorus content had dropped only modestly, but after remelting, almost all the phosphorus had moved from the iron into the slag (Tim Young pers. comm).

Evenstad's eight chapter (translation in Wagner 1990) describes a two-step remelting process in an open hearth very similar to a typical forge. In the the first step, the tuyere is an inch above the hearth floor. The iron bloom is held just above the tuyere until it melts, all the while being fluxed with sand and crushed slag gathered after the initial cutting of the bloom. The remelted iron forms a new bloom below the tuyere. This first melting is the step that removes the phosphorus.

Evenstad then goes on to describe remelting this iron to carburize it into steel, which does not concern us here. His account is marvelously clear compared to most older writings on metallurgy, perhaps because it was a rare instance of the workman and the writer being the same person. To build a remelting hearth such as this is quite a quick and simple job. I build it from either refractory clay, or bricks mortared with clay. I build it with interior dimensions of approximately 25 cm x 30 cm. With clay, it will be oval in plan, or if built with bricks, it will be rectangular. I usually use the same size copper tuyère that I use in the bloomery, though a somewhat smaller inner diameter tuyère, say 1.5 cm instead of 2 cm, works quite well also.

The tuyère enters the narrow side of the hearth approximately 15 cm above the bottom, and protrudes 5 or 6 cm into the hearth. Though Evenstad seems to suggest that the tuyère should be level, I have found that a bit of downward angle, 10 ° to 15 °, helps to center the heat in the hearth, thus keeping your new remelted bloom a bit less crowded against the tuyère, and giving it a bit more room to grow. But as Evenstad suggests, the tuyère should point right along the central axis of the hearth, so that the blast, and thus the heat, is well centered.

With my tuyère set firmly in place, I continue the walls of the hearth up for another 15 cm or so, but leave open the front of the 30 cm side, from which I will be working. This opening should be big enough to stack a couple of bricks in it. These bricks allow me to keep a deep fuel bed in the early stages of the melt, but can be removed when it is time to lift out the resulting "lump" of iron.

If using bricks, I line the walls with a thin layer of clay, perhaps 1.5 to 2 cm thick. Then, of course, I build a small slow fire in the hearth to thoroughly dry the clay. Evenstad does not specify what the floor of the hearth should be made of, just that it should be "pounded flat". During the 2011 experiments, we tried several different things: charcoal fines, sand, crushed ore, but we found that the best floor, by far, was dampened and pounded wood ash. I take wood ash, and dampen it very lightly, mixing and stirring the ash as I dampen it. The ash should not be wet or sticky in any way, just damp enough that it will hold shape when squeezed in the fist, just as a foundryman judges the correct dampness of traditional greensand. With the ash dampened properly, it can be pounded it into a very firm and stable base that can also be scraped out of the way as needed.

Though Evenstad specifies a flat hearth floor 2.5 cm below the tuyère, I prefer a very slightly dished floor, about 4 cm below the bottom of the tuyère orifice. Th dished floor helps to center the hot zone of the fire in the early stages of the melt. The 4 cm depth gives a little more room for the "lump" to grow, while still keeping it in the oxidizing zone of the blast. For the first flux, any silica sand of a reasonably fine particle size will do. For the second iron-rich flux, I do indeed gather the slag from bloom cutting as Evenstad suggests, but I have also successfully used crushed tap slag, hammer scale, or roasted and crushed iron ore. When remelting to remove the phosphorus, I use about twice as much iron-rich flux as sand. I prefer to keep them separate, rather than mixing them, so I can alter the proportions at will.

The iron to be dephosphorized can be in bar form, or raw bloom. I usually use a quarter bloom, weighing 2 kg or so, and that seems to be a convenient size to work with. After the hearth and the bloom are well preheated, the bloom is simply held in front of and above the tuyère until it begins to sparkle and melt on its bottom. Then I begin adding the fluxes, tossing or sprinkling them onto the melting iron, adding them every few minutes until the original bloom has all melted. After the bloom has melted, you may need to push bits of iron and slag into the hot zone to melt better, and of course you may need to clear the tuyère. But as soon as you think you have a nicely formed new "lump", you can pull it out and forge it. And you should now find the phosphorus content of the iron to be much, much lower.

Part three:
The utilization of phosphoric iron

So now we have some methods to influence, if not fully control, our phosphorus content when smelting, and we can remove phosphorus from our resulting bloom, if we want to. But maybe we do not want to keep the phosphorus out. Let us not fall into the trap of thinking about phosphorus as just a contaminant that we should prevent or remove. We should think of it as an alloying material, in the same way we think of carbon. A little bit of either carbon or phosphorus gives us some beneficial properties we can utilize. Just as 2 to 4 % of carbon will give us cast iron, a brittle material of limited uses, so will an excess of phosphorus (what that exact percentage is, I am not trying to address in this paper). But a small amount of carbon in our iron, of say 0.4 to 1 %, gives us one of the most useful materials known to man. So too, does a small amount of phosphorus give us a material with useful properties that purer iron does not have. So why think of phosphorus as a contaminant and carbon as beneficial?

What benefits can we get from the addition of a little phosphorus to our iron? It can improve the forgeability of the iron. It can improve the ease of forge welding. It can significantly harden the iron, especially after cold-working, enough so to make a serviceable edge tool. It slows corrosion when weathering. Let us look at these beneficial properties one at a time.

Forgeability

Being a sculptor, my favorite thing about phosphoric iron is its improved ductility and malleability. It is just a pleasure to forge. This improvement in it is hot-working characteristics is revealed objectively in the hot twist test, but it will be pretty obvious on a more subjective level as soon as you start hitting it, if you are a smith with experience forging different materials. As the phosphorus content of the iron increases, it gets a softer, more buttery feel under the hammer, and the surface takes the impression of the hammer in a very definite and lovely way.

Of course, as with any other nice thing in life, you can overdo it. As the phosphorous content rises into the higher levels that make the iron extremely brittle when cold, it is squishiness when hot becomes extreme, and can lead to a strange little phenomenon. When forging the bloom

Figure 4: Additive sculpture from high phosphorus iron in progress.

into a bar, especially to a square cross section, it will tend to develop longitudinal cracks running along the corners of the bar, which keep reoccurring after welding them. It is as if it is just too slippery inside the bar, and it works back and forth too much as you hammer two sides, cracking and rewelding itself, but leaving either minute cracks or typical signatures of a forge weld. You can see an indication of this in the brittle fracture shown in Figure 9 of the appendix.

Ease of welding

Phosphoric iron has another property of great sculptural benefit: its ease of forge welding. Bloom iron is generally easy to forge weld anyway, but the addition of low levels of phosphorous make it perceptibly easier, welding at lower temperatures. At high levels of phosphorous, again in the range that makes the iron very brittle when cold, it will weld together very easily, at a very low temperature: a yellow color with a hint of orange in it (perhaps 1050 °C). If this iron is brought to a bright white heat, say above 1300 °C, I find that the weld does not stick, and the iron can even crumble in a hot-short manner. But after allowing it to cool slightly, the crumbling will heal, and the weld will stick.

Figure 5: (top) Tamama II, Queen of the Nile. (bottom) The detail shows some of the weld lines from additive construction.

So even an iron with extreme levels of phosphorus has possibilities that can be exploited. For example, in the sculpture pictured here, I used iron that was so high in phosphorous as to be useless for most practical purposes, to create an iron sculpture additively (Fig. 4; Fig. 5). This figure was built up with large wads of very high phosphorous iron. Its softness at forging temperatures and its very low welding temperature allowed me to build it up in exactly the same manner I built its clay model (except the iron needed no armature, of course.) I really do not think I could have done this at all with non-phosphoric iron. Instead of thinking of phosphorus as a contaminant, by simply choosing the proper alloy for the project at hand, I was able to broaden my sculptural possibilities.

Balasubramaniam has argued that the Iron Pillar of Delhi was intentionally made with phosphoric iron for its anti-corrosion properties. But I wonder if perhaps phosphoric iron was also chosen for such large forgings to facilitate the welding (Balasubramaniam 2002, 7).

Figure 6: This knife was forged from low-carbon, medium-phosphorus iron, and the edge was cold hammered before sharpening. I have used it to gut, skin, and butcher several deer.

A small aside on heat treatment

In high-phosphorus, brittle iron, part of the embrittlement is caused by the phosphorus migrating to the grain boundaries. This migration occurs upon cooling, and especially when cooling slowly through the range of 400 to 600 °C (Stewart et al. 2000, 275). So the brittleness can be reduced by quenching the iron from red heat, thus passing through the 200 to 600 degree range as quickly as possible. This seems very backwards to a modern blacksmith used to carbon steels! This strengthening is moderate though. We did a simple set of experiments with the very high-phosphorous brittle iron, using small bars forged from the same part of the same bloom. One bar was air-cooled, and one bar was quenched in water at a red heat We tested the relative strength of the bars by dropping a weight, from increasing 1 inch increment heights, recording the height at which the bar snapped. The heat treatment consistently strengthened the iron, increasing the height of the dropped weight before snapping the bar by about 30 %, on average.

Hardness

The other beneficial properties of phosphoric iron are a little more utilitarian than making a sculptor happy. Just as with carbon, the addition of small amounts of phosphorus makes iron appreciably harder. And that hardness can be increased further by cold hammering, enough so to make a very serviceable blade (Fig. 6). Of course this is well known among archaeometallurgists, but perhaps a few more observations from me will be of some small value. A phosphoric iron knife certainly doesn't take as keen an edge as a heat-treated steel knife, and it does not hold that edge as long. But it is perfectly serviceable, and it is edge can be quickly restored with just a few swipes on a stone.

Figure 7: Hammer of low-carbon phosphoric iron, showing only mild mushrooming after many years of use (hot forging only).

Even a moderate increase in hardness can be quite useful in tools that are less reliant on a keen edge. For example, think of a hoe, which one could argue is the most important tool we've invented so far. A phosphoric iron hoe, being harder than a pure wrought iron one, will last longer, since it will not abrade as quickly. Indeed, the act of using it will work harden the tool to some extent while using it, and even a farmer can hammer on the edge to harden it some more. A forging hammer is another excellent use of phosphoric iron (Fig. 7): the hammer will mostly be hitting hot iron, so does not need to be terribly hard, and the act of using it work-hardens the face.

The scythe is a tool that is still hammer-hardened and honed by the user. A scytheman carries a small anvil and hammer into the field, which he uses to reshape and reharden the edge during the day. Though scythes in more recent centuries have been made of various steel and iron combinations, I wonder if the roots of the technique arose from using phosphoric iron? I did a fairly comprehensive set of hardness tests on samples of my phosphoric iron from many different smelts. I measured these using Rockwell scale A (HRA). I used this scale because it was most appropriate to the range of hardness I was examining, and was available to me in a well-calibrated machine at a nearby university. But here I will give the hardness in the Vickers scale, though I know these sorts of translations are only a rough correlation. All the values I give here represent an average of many measurements across the sample.

For comparisons sake, a piece of "Pure Iron" (a commercially available very low carbon iron) measured 103 HV. A sample of hot roll mild steel measured 112 HV.

Hot worked and air-cooled phosphoric iron ranged in hardness from 75 to 156 HV, but most of the good workable iron was in the range of 118 to 141 HV. The measurements of 145 to 156 HV were only on extremely high-phosphorus, brittle iron that I forged into knives to see how they would serve. This iron was too brittle to be able to harden the edge by cold-hammering. Hot forged and aircooled, they were hard enough to take a somewhat decent edge, but they were too brittle for the blades to survive hard service.

Most interesting is the variation of hardness between hot-worked and cold worked iron of more moderate phosphorus content. Three hot forged bars, all with a starting hardness close to 141 HV, were reduced about 40 % in thickness by cold hammering. These bars then measured from 188 to 230 HV. Other researchers have reported even higher increases in hardness, up to 340 HV (Tylecote 1986, 145). Many archaeometallurgical reports show that this degree of hardness was considered perfectly adequate for knives, and even swords, even into the medieval period (Tylecote 1986 173-174; Lang & Williams 1975, 199; Wilthew 1987, 62-74).

Corrosion Resistance

Bloom iron, phosphoric or no, is already more corrosion resistant than modern mild steel. There is less danger of sulfur from the fuel with charcoal than coke (though I have indeed made hot-short iron in which the sulfur came from the charcoal.) In modern steels, sulfur is removed by the addition of metallic manganese, which leaves behind small particles of manganese sulfide, which are known points of incipient corrosion, which leads to the characteristic “pitting” of modern steels. Bloom iron avoids these problems.

It is of course well known that phosphorus can improve this already decent corrosion resistance. When the iron begins to rust, the phosphorus forms iron phosphates, creating a protective film, especially at the rust/metal interface (Balasubramaniam 2002, 89-119). But again, I simply offer an observation or two about corrosion and patina of phosphoric iron from a smith’s perspective.

You might not find it surprising to learn that I have quite a few unsold sculptures living outdoors around my place, and so I have been able to observe their rusting, or lack thereof, for many years now. The phosphoric iron, besides being much slower to rust in general, tends to build that rust into a thin, dark brown layer that I find quite attractive. The steels or antique wrought pieces that are their neighbors tend to rust to slightly yellower browns, rust more quickly, and build thicker, flakier layers of rust.

One last property of phosphoric iron that we should briefly discuss is its use for contrast in pattern welding. In discussions with bladesmiths, both in person and online, they have often referred to phosphorus iron being used as the bright lines in Anglo-Saxon and Viking age pattern-welded swords. My own experience with pattern-welding phosphoric iron has been a bit more unclear and confusing, and makes me doubt the issue is quite as simple as that. In using both low and medium phosphorus iron to contrast with another iron or steel, I have sometimes found them to etch dark in comparison, at least in the etchants I have tried (nital, ferric chloride, and vinegar, to name a few). I have other colleagues that have observed the same thing.

On the other hand, I forged a laminate iron, of high-phosphorus and medium- phosphorus iron, smelted from the same phosphoric ore but under the two different smelting regimes. I welded it into a 6 layer billet and forged it down. This did give a bar with nice contrast after etching. I took this to the hardness tester, and indeed the bright lines had a higher average hardness (average 155HV) than the dark lines (average 100 HV). But the bright line issue might be confused by the fact that the phosphorus does seem to migrate to the welds, and so does nickel when present, and so does arsenic, which is sometimes present in the ores and often present in forging coal. So I think the question of the etching behavior of phosphoric iron in pattern-welding is unclear, and perhaps bears further experimental study.

Conclusion

So I have presented here not proofs, but indications, that the level of phosphorus in bloom iron can be manipulated by smelting procedures. Phosphorus uptake can be inhibited by a moderate temperature and smelting rate, moderate reducing conditions, careful management of slag tapping, and limiting the size and yield of the bloom. I have demonstrated that phosphorus can be removed from iron by remelting in a simple hearth. And I have discussed the benefits of phosphorus in iron: forgeability, weldability, improved hardness, and corrosion resistance. I do not pretend to a complete understanding of the behavior of phosphorus in bloom iron. But I hope my perspectives on these issues can at least assist and inspire further work by other smelters, smiths, and archaeologists. There is only one thing I am certain about here: that the presence of phosphorus in your iron should not be viewed as a contaminant to be avoided, but a set of opportunities to be embraced.

Acknowledgements

To thank the people who have assisted and inspired me in this work would literally run on for pages. So I will just have to settle for thanking the two people without whom it would have actually been impossible: Elizabeth Sauder and Skip Williams.

Appendix: Twist tests

Here is our standard method for both hot and cold twist tests. For the hot twist test, I forge a section of bloom to 16 mm (5/8 inch) square bar (Fig. 8). Of course the section of bloom chosen to forge is a major variable here. In general practice, I section a bloom pie-wise into halves, quarters, or eighths as desired, and for the test, forge a bar from the more solid iron at the center of the bloom. This bar is forged directly from the bloom, without folding and welding several layers, as many people think is necessary. This direct forging to bar is itself a first test: not of the phosphorus content, but of the general quality of your bloom. If you cannot get a decent bar without folding and welding, you still have work to do to improve your smelting, regardless of phosphorus content.

With a center punch, I mark off a 75 mm (3 inch) long section of this bar, and heat it to a nice butter yellow. Quickly clamping it in the vise at one punch mark, and grabbing it with the twisting wrench at the other, I twist it until it shears in two, counting each quarter turn as I go. Because of the vagaries of heating, this test is less precise than the cold twist test, but it does give some excellent subjective information. Of course, there are many factors that affect how a bar survives the hot twist test, including the density of the bloom, its relative homogeneity, and your skill at forging it to bar. But as for its applicability to phosphorus content, low-phosphorus bars will usually begin to tear apart between 3 and 4 complete twists, and higher-phosphorus bars will go four to five twists before failing, due to the high-phosphorus iron being more ductile at forging temperatures.

Perhaps it is possible to have a high enough phosphorus content to make the metal too hot-short to forge at all, which has been reported to me by other experimenters. But I have never experienced this, myself. Occasionally in a very high-phosphorus, brittle iron, I have experienced hot-short crumbling at a full white heat, but this material then heals and welds upon cooling to yellow. More often though, the higher-phosphorus metal forges and hot twists so beautifully that it is hard to imagine


Figure 8: Bar forged direct from bloom.


Figure 9: High-phosphorus brittle fracture. You can also see here the typical diagonal cracks from the corners discussed in part three.


Figure 10: Low-phosphorus ductile fracture.

anything could possibly be amiss with it, until you get to the cold twist test.

For the cold twist test, I forge a bar down to 8 mm (5/16 inch) square bar, finishing the forging at a bright red heat, and then allow the bar to air cool. I mark off a 38 mm (1.5 inch) long section of bar, clamp it in the vise at one mark, and once again twist the section until it breaks. With this test, you not only get a count of the twists, but the fracture surface thus exposed is extremely revealing. I am able to generally perceive three levels of phosphorus by this test.

High-phosphorus bars will give with a sudden snap, surviving zero to half of a revolution. The fracture will show large and brilliant grains, across the entire break (Fig. 9). Low-phosphorus bars will survive 1 revolution or better, and show a ductile, smeary fracture, with a grayer, duller surface (Fig. 10). The best result I have achieved to date in any bloom iron, even with phosphorus-free ore, is two full twists before failure. Bars of a moderate phosphorus content will go from half to one twist before failure, and show a mix of ductile and crystalline fractures. What the actual quantitative levels of phosphorus these three grades of iron represent will have to await future experiment and analysis.

Bibliography

Balasubramaniam, R. (2002). *Delhi Iron Pillar: new insights*. New Delhi: Aryan Books International.

Dillmann, P., & Balasubramaniam, R. (2001). Characterization of ancient Indian iron and entrapped slag inclusions using electron, photon and nuclear microprobes. *Bulletin of Materials Science* 24, 317–322. doi: 10.1007/BF02704929

Godfrey, E. G. (2007). *The technology of ancient and medieval directly reduced phosphoric iron* (PhD thesis, University of Bradford).

Kumar, V. & Balasubramanium, R. (2002). On the origin of high phosphorus contents in ancient Indian iron. *Metals, Materials and Process* 14(1), 1-14. doi: 10.1016/S1471-3918(02)80023-5

Lang, J. & Williams, A. R. (1975). The hardening of iron swords. *Journal of Archaeological Science* 2, 199–207. doi: 10.1016/0305-4403(75)90057-6

Sauder, L. (2013a). An American bloomery in Sussex. In: D. Dungworth & R. C. P. Doonan (Eds.), *Accidental and experimental archaeometallurgy* (pp. 69–74). London: The Historical Metallurgy Society (HMS Occasional Publication 7).

Sauder, L. (2013b). Making steel in the "Aristotle furnace'. In: D. Dungworth & R. C. P. Doonan (Eds.), *Accidental and experimental archaeometallurgy* (pp. 107–110). London: The Historical Metallurgy Society (HMS Occasional Publication 7).

Stepanov, I., Sauder, L., Keen, J., Workman, V. & Eliyahu-Behar, A. (2022). By the hand of the smelter: tracing the impact of decision-making in bloomery iron smelting. *Archaeological and Anthropological Sciences* 14, 80. doi: 10.1007/s12520-022-01516-3

Stewart, J., Charles, J. & Wallach, E. (2000). Iron-phosphorus-carbon system Part 1. Mechanical properties of low carbon iron-phosphorus alloys. *Materials Science and Technology* 16, 275–282. doi: 10.1179/026708300101507839

Tylecote R. F. (1986). *The Prehistory of metallurgy in the British Isles*. London: Institute of Metals.

Vike, V. (2016). *Refining iron by re-smelting*. Retrieved December 12, 2024, from: https://www.khm.uio.no/english/research/projects/langeid/project/iron-steel/resmelt-iron/index.html

Vizcaino, A., McColm, I. J., & McDonnell, J. G. (1997). Phosphorus in the Bloomery Iron Process. *Caesaraugusta* 73, 237–243.

Wagner, D. (1990). Ancient carburization of iron to steel: a comment. *Archeomaterials* 4(1), 111–117.

Wilthew, P. (1987). Metallographic examination of medieval knives and shears. In: J. Cowgill, M. De Neergaard & N. Griffiths (Eds.), *Knives and scabbards. Medieval finds from excavations in London* (pp. 62–74). London: Museum of London.

The Authors

Lee Sauder is a professional blacksmith, sculptor, and bloom smelter. He began blacksmithing in 1973 at the tender age of 12, and had never stopped since. He lives and works on the south slope of House Mountain in Rockbridge County, Virginia.
Address: 1365 Bird Forest Rd, Lexington VA 24450, USA
Email: sauderironwork@gmail.com

Originally trained as an electrical engineer, Skip Williams is a polymath and independent scholar whose published works include botanical genetics as well as experimental archaeometallurgy. His home, shop and lab are in a log cabin just over the ridge from Lee's shop.
Address: 1038 Sugar Creek Rd, Lexington VA 24450, USA
Email: skip.williams@yahoo.com
ORCID: https://orcid.org/0000-0003-3920-8565